

The Tree Width of Separation Logic with Recursive Definitions

Radu Iosif¹, Adam Rogalewicz², and Jiri Simacek²

¹ VERIMAG/CNRS, Grenoble, France

² FIT, Brno University of Technology, IT4Innovations Centre of Excellence, Czech Republic

Abstract. Separation Logic is a widely used formalism for describing dynamically allocated linked data structures, such as lists, trees, etc. The decidability status of various fragments of the logic constitutes a long standing open problem. Current results report on techniques to decide satisfiability and validity of entailments for Separation Logic(s) over lists (possibly with data). In this paper we establish a more general decidability result. We prove that any Separation Logic formula using rather general recursively defined predicates is decidable for satisfiability, and moreover, entailments between such formulae are decidable for validity. These predicates are general enough to define (doubly-) linked lists, trees, and structures more general than trees, such as trees whose leaves are chained in a list. The decidability proofs are by reduction to decidability of Monadic Second Order Logic on graphs with bounded tree width.

1 Introduction

Separation Logic (SL) [17] is a general framework for describing dynamically allocated mutable data structures generated by programs that use pointers and low-level memory allocation primitives. The logics in this framework are used by an important number of academic (SPACE INVADER [1], SLEEK [16] and PREDATOR [9]), as well as industrial-scale (INFER [7]) tools for program verification and certification. These logics are used both externally, as property specification languages, or internally, as e.g., abstract domains for computing invariants, or for proving verification conditions. The main advantage of using SL when dealing with heap manipulating programs, is the ability to provide compositional proofs, based on the principle of *local reasoning* i.e., analyzing different sections (e.g., functions, threads, etc.) of the program, that work on disjoint parts of the global heap, and combining the analysis results a-posteriori.

The basic language of SL consists of two kinds of atomic propositions describing either (i) the empty heap, or (ii) a heap consisting of an allocated cell, connected via a separating conjunction primitive. Hence a basic SL formula can describe only a heap whose size is bounded by the size of the formula. The ability of describing unbounded data structures is provided by the use of *recursive definitions*. Figure 1 gives several common examples of recursive data structures definable in this framework.

The main difficulty that arises when using Separation Logic with Recursive Definitions (SLRD) to reason automatically about programs is that the logic, due to its expressiveness, does not have very nice decidability properties. Most dialects used in practice restrict the language (e.g., no quantifier alternation, the negation is used in a

$$\begin{aligned}
list(hd, tl) &::= emp \wedge hd = tl \\
&\quad | \quad \exists x. hd \mapsto x * list(x, tl) \\
dll(hd, p, tl) &::= emp \wedge hd = tl \\
&\quad | \quad \exists x. hd \mapsto (x, p) * dll(x, hd, tl) \\
tree(root) &::= emp \wedge root = nil \\
&\quad | \quad \exists l, r. root \mapsto (l, r) * tree(l) * tree(r) \\
tll(x, ll, lr) &::= x \mapsto (nil, nil, lr) \wedge x = ll \\
&\quad | \quad \exists l, r, z. x \mapsto (l, r, nil) * tll(l, ll, z) \\
&\quad \quad * tll(r, z, lr)
\end{aligned}$$

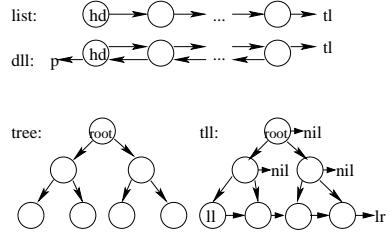


Fig. 1. Examples of recursive data structures definable in SLRD.

very restricted ways, etc.) and the class of models over which the logic is interpreted (typically singly-linked lists, and slight variations thereof). In the same way, we apply several natural restrictions on the syntax of the recursive definitions, and define the fragment SLRD_{btw} , which guarantees that all models of a formula in the fragment have *bounded tree width*. Indeed, this ensures that the satisfiability and entailment problems in this fragment are decidable *without any restrictions on the type of the recursive data structures considered*.

In general, the techniques used in proving decidability of Separation Logic are either proof-based ([16, 2]), or model-based ([5, 8]). It is well-known that automata theory, through various automata-logics connections, provides a unifying framework for proving decidability of various logics, such as (W)Sks, Presburger Arithmetic or MSO over certain classes of graphs. In this paper we propose an automata-theoretic approach consisting of two ingredients. First, SLRD_{btw} formulae are translated into equivalent Monadic Second Order (MSO) formulae over graphs. Second, we show that the models of SLRD_{btw} formulae have the *bounded tree width* property, which provides a decidability result by reduction to the satisfiability problem for MSO interpreted over graphs of bounded tree width [18], and ultimately, to the emptiness problem of tree automata.

Related Work The literature on defining decidable logics for describing mutable data structures is rather extensive. Initially, first-order logic with transitive closure of one function symbol was introduced in [11] with a follow-up logic of reachability on complex data structures, in [19]. The decision procedures for these logics are based on reductions to the decidability of MSO over finite trees. Along the same lines, the logic PALE [15] goes beyond trees, in defining trees with edges described by regular routing expressions, whose decidability is still a consequence of the decidability of MSO over trees. More recently, the CSL logic [4] uses first-order logic with reachability (along multiple selectors) in combination with arithmetic theories to reason about shape, path lengths and data within heap structures. Their decidability proof is based on a small model property, and the algorithm is enumerative. In the same spirit, the STRAND logic [14] combines MSO over graphs, with quantified data theories, and provides decidable fragments using a reduction to MSO over graphs of bounded tree width.

On what concerns SLRD [17], the first (proof-theoretic) decidability result on a restricted fragment defining only singly-linked lists was reported in [2], which describe a coNP algorithm. The full basic SL without recursive definitions, but with the magic

wand operator was found to be undecidable when interpreted *in any memory model* [6]. Recently, the entailment problem for SLRD over lists has been reduced to graph homomorphism in [8], and can be solved in PTIME. This method has been extended to reason nested and overlaid lists in [10]. The logic SLRD_{btw}, presented in this paper is, to the best of our knowledge, the first decidable SL that can define structures more general than lists and trees, such as e.g. trees with parent pointers and linked leaves.

2 Preliminaries

For a finite set S , we denote by $\|S\|$ its cardinality. We sometimes denote sets and sequences of variables as \mathbf{x} , the distinction being clear from the context. If \mathbf{x} denotes a sequence, $(\mathbf{x})_i$ denotes its i -th element. For a partial function $f : A \rightharpoonup B$, and $\perp \notin B$, we denote $f(x) = \perp$ the fact that f is undefined at some point $x \in A$. By $f[a \leftarrow b]$ we denote the function $\lambda x . \text{if } x = a \text{ then } b \text{ else } f(x)$. The domain of f is denoted $\text{dom}(f) = \{x \in A \mid f(x) \neq \perp\}$, and the image of f is denoted as $\text{img}(f) = \{y \in B \mid \exists x \in A . f(x) = y\}$. By $f : A \rightharpoonup_{\text{fin}} B$ we denote any partial function whose domain is finite. Given two partial functions f, g defined on disjoint domains, we denote by $f \oplus g$ their union.

Stores, Heaps and States. We consider $PVar = \{u, v, w, \dots\}$ to be a countable infinite set of *pointer variables* and $Loc = \{l, m, n, \dots\}$ to be a countable infinite set of *memory locations*. Let $\text{nil} \in PVar$ be a designated variable, $\text{null} \in Loc$ be a designated location, and $\text{Sel} = \{1, \dots, S\}$, for some given $S > 0$, be a finite set of natural numbers, called *selectors* in the following.

Definition 1. A state is a pair $\langle s, h \rangle$ where $s : PVar \rightharpoonup Loc$ is a partial function mapping pointer variables into locations such that $s(\text{nil}) = \text{null}$, and $h : Loc \rightharpoonup_{\text{fin}} Sel \rightharpoonup_{\text{fin}} Loc$ is a finite partial function such that (i) $\text{null} \notin \text{dom}(h)$ and (ii) for all $\ell \in \text{dom}(h)$ there exist $k \in \text{Sel}$ such that $(h(\ell))(k) \neq \perp$.

Given a state $S = \langle s, h \rangle$, s is called the *store* and h the *heap*. For any $k \in \text{Sel}$, we write $h_k(\ell)$ instead of $(h(\ell))(k)$, and $\ell \xrightarrow{k} \ell'$ for $h_k(\ell) = \ell'$. We sometimes call a triple $\ell \xrightarrow{k} \ell'$ an *edge*, and k is called a *selector*. Let $\text{Img}(h) = \bigcup_{\ell \in Loc} \text{img}(h(\ell))$ be the set of locations which are destinations of some selector edge in h . A location $\ell \in Loc$ is said to be *allocated* in $\langle s, h \rangle$ if $\ell \in \text{dom}(h)$ (i.e. it is the source of an edge), and *dangling* in $\langle s, h \rangle$ if $\ell \in [\text{img}(s) \cup \text{Img}(h)] \setminus \text{dom}(h)$, i.e., it is either referenced by a store variable, or reachable from an allocated location in the heap, but it is not allocated in the heap itself. The set $\text{loc}(S) = \text{img}(s) \cup \text{dom}(h) \cup \text{Img}(h)$ is the set of all locations either allocated or referenced in a state $S = \langle s, h \rangle$.

Trees. Let Σ be a finite label alphabet, and \mathbb{N}^* be the set of sequences of natural numbers. Let $\epsilon \in \mathbb{N}^*$ denote the empty sequence, and $p.q$ denote the concatenation of two sequences $p, q \in \mathbb{N}^*$. A *tree* t over Σ is a finite partial function $t : \mathbb{N}^* \rightharpoonup_{\text{fin}} \Sigma$, such that $\text{dom}(t)$ is a finite prefix-closed subset of \mathbb{N}^* , and for each $p \in \text{dom}(t)$ and $i \in \mathbb{N}$, we have: $t(p.i) \neq \perp \Rightarrow \forall 0 \leq j < i . t(p.j) \neq \perp$. Given two positions $p, q \in \text{dom}(t)$, we say

that q is the i -th successor (child) of p if $q = p.i$, for $i \in \mathbb{N}$. Also q is a successor of p , or equivalently, p is the parent of q , denoted $p = \text{parent}(q)$ if $q = p.i$, for some $i \in \mathbb{N}$.

We will sometimes denote by $\mathcal{D}(t) = \{-1, 0, \dots, N\}$ the *direction alphabet* of t , where $N = \max\{i \in \mathbb{N} \mid p.i \in \text{dom}(t)\}$. The concatenation of positions is defined over $\mathcal{D}(t)$ with the convention that $p.(-1) = q$ if and only if $p = q.i$ for some $i \in \mathbb{N}$. We denote $\mathcal{D}_+(t) = \mathcal{D}(t) \setminus \{-1\}$. A *path* in t , from p_1 to p_k , is a sequence $p_1, p_2, \dots, p_k \in \text{dom}(t)$ of pairwise distinct positions, such that either $p_i = \text{parent}(p_{i+1})$ or $p_{i+1} = \text{parent}(p_i)$, for all $1 \leq i < k$. Notice that a path in the tree can also link sibling nodes, not just ancestors to their descendants, or viceversa. However, a path may not visit the same tree position twice.

Tree Width. A state (Def. 1) can be seen as a directed graph, whose nodes are locations, and whose edges are defined by the selector relation. Some nodes are labeled by program variables (*PVar*) and all edges are labeled by selectors (*Sel*). The notion of tree width is then easily adapted from generic labeled graphs to states. Intuitively, the tree width of a state (graph) measures the similarity of the state to a tree.

Definition 2. Let $S = \langle s, h \rangle$ be a state. A tree decomposition of S is a tree $t : \mathbb{N}^* \rightharpoonup_{\text{fin}} 2^{\text{loc}(S)}$, labeled with sets of locations from $\text{loc}(S)$, with the following properties:

1. $\text{loc}(S) = \bigcup_{p \in \text{dom}(t)} t(p)$, the tree covers the locations of S
2. for each edge $l_1 \xrightarrow{s} l_2$ in S , there exists $p \in \text{dom}(t)$ such that $l_1, l_2 \in t(p)$
3. for each $p, q, r \in \text{dom}(t)$, if q is on a path from p to r in t , then $t(p) \cap t(r) \subseteq t(q)$

The width of the decomposition is $w(t) = \max_{p \in \text{dom}(t)} \{\|t(p)\| - 1\}$. The tree width of S is $\text{tw}(S) = \min\{w(t) \mid t \text{ is a tree decomposition of } S\}$.

A set of states is said to have *bounded tree width* if there exists a constant $k \geq 0$ such that $\text{tw}(S) \leq k$, for any state S in the set. Figure 2 gives an example of a graph (left) and a possible tree decomposition (right).

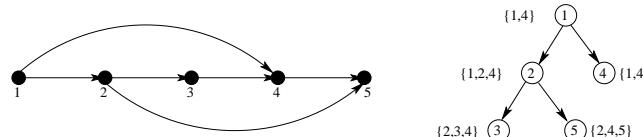


Fig. 2. A graph and a possible tree decomposition of width 2

2.1 Syntax and Semantics of Monadic Second Order Logic

Monadic second-order logic (MSO) on states is a straightforward adaptation of MSO on labeled graphs [13]. As usual, we denote first-order variables, ranging over locations,

by x, y, \dots , and second-order variables, ranging over sets of locations, by X, Y, \dots . The set of logical MSO variables is denoted by $LVar_{mso}$, where $PVar \cap LVar_{mso} = \emptyset$.

We emphasize here the distinction between the logical variables $LVar_{mso}$ and the pointer variables $PVar$: the former may occur within the scope of first and second order quantifiers, whereas the latter play the role of symbolic constants (function symbols of zero arity). For the rest of this paper, a logical variable is said to be free if it does not occur within the scope of a quantifier. By writing $\varphi(\mathbf{x})$, for an MSO formula φ , and a set of logical variables \mathbf{x} , we mean that all free variables of φ are in \mathbf{x} .

The syntax of MSO is defined below:

$$\begin{aligned} u &\in PVar; x, X \in LVar_{mso}; k \in \mathbb{N} \\ \varphi ::= &x = y \mid var_u(x) \mid edge_k(x, y) \mid null(x) \mid X(x) \mid \varphi \wedge \varphi \mid \neg \varphi \mid \exists x. \varphi \mid \exists X. \varphi \end{aligned}$$

The semantics of MSO on states is given by the relation $S, \iota, v \models_{mso} \varphi$, where $S = \langle s, h \rangle$ is a state, $\iota : \{x, y, z, \dots\} \rightarrow_{fin} Loc$ is an interpretation of the first order variables, and $v : \{X, Y, Z, \dots\} \rightarrow_{fin} 2^{Loc}$ is an interpretation of the second order variables. If $S, \iota, v \models_{mso} \varphi$ for all interpretations $\iota : \{x, y, z, \dots\} \rightarrow_{fin} Loc$ and $v : \{X, Y, Z, \dots\} \rightarrow_{fin} 2^{Loc}$, then we say that S is a *model* of φ , denoted $S \models_{mso} \varphi$. We use the standard MSO semantics [18], with the following interpretations of the vertex and edge labels:

$$\begin{aligned} S, \iota, v \models_{mso} null(x) &\iff \iota(x) = nil \\ S, \iota, v \models_{mso} var_u(x) &\iff s(u) = \iota(x) \\ S, \iota, v \models_{mso} edge_k(x, y) &\iff h_k(\iota(x)) = \iota(y) \end{aligned}$$

The *satisfiability problem* for MSO asks, given a formula φ , whether there exists a state S such that $S \models_{mso} \varphi$. This problem is, in general, undecidable. However, one can show its decidability on a restricted class of models. The theorem below is a slight variation of a classical result in (MSO-definable) graph theory [18]. For space reasons, all proofs are given in [12].

Theorem 1. *Let $k \geq 0$ be an integer constant, and φ be an MSO formula. The problem asking if there exists a state S such that $tw(S) \leq k$ and $S \models_{mso} \varphi$ is decidable.*

2.2 Syntax and Semantics of Separation Logic

Separation Logic (SL) [17] uses only a set of first order logical variables, denoted as $LVar_{sl}$, ranging over locations. We suppose that $LVar_{sl} \cap PVar = \emptyset$ and $LVar_{sl} \cap LVar_{mso} = \emptyset$. Let Var_{sl} denote the set $PVar \cup LVar_{sl}$. A formula is said to be *closed* if it does not contain logical variables which are not under the scope of a quantifier. By writing $\varphi(\mathbf{x})$ for an SL formula φ and a set of logical variables \mathbf{x} , we mean that all free variables of φ are in \mathbf{x} .

Basic Formulae. The syntax of basic formula is given below:

$$\begin{aligned} \alpha &\in Var_{sl} \setminus \{nil\}; \beta \in Var_{sl}; x \in LVar_{sl} \\ \pi ::= &\alpha = \beta \mid \alpha \neq \beta \mid \pi_1 \wedge \pi_2 \\ \sigma ::= &emp \mid \alpha \mapsto (\beta_1, \dots, \beta_n) \mid \sigma_1 * \sigma_2, \text{ for some } n > 0 \\ \varphi ::= &\pi \wedge \sigma \mid \exists x. \varphi \end{aligned}$$

A formula of the form $\bigwedge_{i=1}^n \alpha_i = \beta_i \wedge \bigwedge_{j=1}^m \alpha_j \neq \beta_j$ defined by π in the syntax above is said to be *pure*. If Π is a pure formula, let Π^* denote its *closure*, i.e., the equivalent pure formula obtained by the exhaustive application of the reflexivity, symmetry, and transitivity axioms of equality. A formula of the form $\star_{i=1}^k \alpha_i \mapsto (\beta_{i,1}, \dots, \beta_{i,n})$ defined by σ in the syntax above is said to be *spatial*. The atomic proposition *emp* denotes the empty spatial conjunction. For a spatial formula Σ , let $|\Sigma|$ be the total number of variable occurrences in Σ , e.g. $|\text{emp}| = 0$, $|\alpha \mapsto (\beta_1, \dots, \beta_n)| = n + 1$, etc.

The semantics of a basic formula φ is given by the relation $S, \iota \models_{sl} \varphi$ where $S = \langle s, h \rangle$ is a state, and $\iota : LVar_{sl} \rightarrow_{fin} Loc$ is an interpretation of logical variables from φ . For a closed formula φ , we denote by $S \models_{sl} \varphi$ the fact that S is a *model* of φ .

$$\begin{aligned} S, \iota \models_{sl} \text{emp} &\iff \text{dom}(h) = \emptyset \\ S, \iota \models_{sl} \alpha \mapsto (\beta_1, \dots, \beta_n) &\iff h = \{\langle (s \oplus \iota)(\alpha), \lambda i . \text{if } i \leq n \text{ then } (s \oplus \iota)(\beta_i) \text{ else } \perp \}\} \\ S, \iota \models_{sl} \varphi_1 * \varphi_2 &\iff S_1, \iota \models_{sl} \varphi_1 \text{ and } S_2, \iota \models_{sl} \varphi_2 \text{ where } S_1 \uplus S_2 = S \end{aligned}$$

The semantics of $=$, \neq , \wedge , and \exists is classical. Here, the notation $S_1 \uplus S_2 = S$ means that S is the union of two states $S_1 = \langle s_1, h_1 \rangle$ and $S_2 = \langle s_2, h_2 \rangle$ whose stacks agree on the evaluation of common program variables ($\forall \alpha \in PVar . s_1(\alpha) \neq \perp \wedge s_2(\alpha) \neq \perp \Rightarrow s_1(\alpha) = s_2(\alpha)$), and whose heaps have disjoint domains ($\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$) i.e., $S = \langle s_1 \cup s_2, h_1 \oplus h_2 \rangle$. Note that we adopt here the *strict semantics*, in which a points-to relation $\alpha \mapsto (\beta_1, \dots, \beta_n)$ holds in a state consisting of a single cell pointed to by α , with exactly n outgoing edges towards dangling locations pointed to by β_1, \dots, β_n , and the empty heap is specified by *emp*.

Every basic formula φ is equivalent to an existentially quantified pair $\Sigma \wedge \Pi$ where Σ is a spatial formula and Π is a pure formula. Given a basic formula φ , one can define its spatial (Σ) and pure (Π) parts uniquely, up to equivalence. A variable $\alpha \in Var$ is said to be *allocated in* φ if and only if $\alpha \mapsto (\dots)$ occurs in Σ . It is easy to check that an allocated variable may not refer to a dangling location in any model of φ . A variable β is *referenced* if and only if $\alpha \mapsto (\dots, \beta, \dots)$ occurs in Σ for some variable α . For a basic formula $\varphi \equiv \Sigma \wedge \Pi$, the *size* of φ is defined as $|\varphi| = |\Sigma|$.

Lemma 1. *Let $\varphi(\mathbf{x})$ be a basic SL formula, $S = \langle s, h \rangle$ be a state, and $\iota : LVar_{sl} \rightarrow_{fin} Loc$ be an interpretation, such that $S, \iota \models_{sl} \varphi(\mathbf{x})$. Then $tw(S) \leq \max(|\varphi|, \|PVar\|)$.*

Recursive Definitions. A system \mathcal{P} of recursive definitions is of the form:

$$\begin{aligned} P_1(x_{1,1}, \dots, x_{1,n_1}) &::= |_{j=1}^{m_1} R_{1,j}(x_{1,1}, \dots, x_{1,n_1}) \\ &\dots \\ P_k(x_{k,1}, \dots, x_{k,n_k}) &::= |_{j=1}^{m_k} R_{k,j}(x_{k,1}, \dots, x_{k,n_k}) \end{aligned}$$

where P_1, \dots, P_k are called *predicates*, $x_{i,1}, \dots, x_{i,n_i}$ are called *parameters*, and the formulae $R_{i,j}$ are called the *rules* of P_i . Concretely, a rule $R_{i,j}$ is of the form $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} . \Sigma * P_{i_1}(\mathbf{y}_1) * \dots * P_{i_m}(\mathbf{y}_m) \wedge \Pi$, where Σ is a spatial SL formula over variables $\mathbf{x} \cup \mathbf{z}$, called the *head* of $R_{i,j}$, $\langle P_{i_1}(\mathbf{y}_1), \dots, P_{i_m}(\mathbf{y}_m) \rangle$ is an ordered sequence of *predicate occurrences*, called the *tail* of $R_{i,j}$ (we assume w.l.o.g. that $\mathbf{x} \cap \mathbf{z} = \emptyset$, and that $\mathbf{y}_k \subseteq \mathbf{x} \cup \mathbf{z}$, for all $k = 1, \dots, m$), Π is a pure formula over variables $\mathbf{x} \cup \mathbf{z}$.

Without losing generality, we assume that all variables occurring in a rule of a recursive definition system are logical variables from $LVar_{sl}$ – pointer variables can be passed as parameters at the top level. We subsequently denote $head(R_{i,j}) \equiv \Sigma$, $tail(R_{i,j}) \equiv \langle P_{i_k}(\mathbf{y}_k) \rangle_{k=1}^m$ and $pure(R_{i,j}) \equiv \Pi$, for each rule $R_{i,j}$. Rules with empty tail are called *base cases*. For each rule $R_{i,j}$ let $\|R_{i,j}\|^{var} = \|\mathbf{z}\| + \|\mathbf{x}\|$ be the number of variables, both existentially quantified and parameters, that occur in $R_{i,j}$. We denote by $\|\mathcal{P}\|^{var} = \max\{\|R_{i,j}\|^{var} \mid 1 \leq i \leq k, 1 \leq j \leq m_i\}$ the maximum such number, among all rules in \mathcal{P} . We also denote by $\mathcal{D}(\mathcal{P}) = \{-1, 0, \dots, \max\{|tail(R_{i,j})| \mid 1 \leq i \leq k, 1 \leq j \leq m_i\} - 1\}$ the *direction alphabet* of \mathcal{P} .

Example. The predicate tll describes a data structure called a *tree with parent pointers and linked leaves* (see Fig. 3(b)). The data structure is composed of a binary tree in which each internal node points to left and right children, and also to its parent node. In addition, the leaves of the tree are kept in a singly-linked list, according to the order in which they appear on the frontier (left to right).

$$\begin{aligned} tll(x, p, leaf_l, leaf_r) ::= & x \mapsto (nil, nil, p, leaf_r) \wedge x = leaf_l & (R_1) \\ | & \exists l, r, z. x \mapsto (l, r, p, nil) * tll(l, x, leaf_l, z) * tll(r, x, z, leaf_r) & (R_2) \end{aligned}$$

The base case rule (R_1) allocates leaf nodes. The internal nodes of the tree are allocated by the rule (R_2) , where the tll predicate occurs twice, first for the left subtree, and second for the right subtree. \square

Definition 3. Given a system of recursive definitions $\mathcal{P} = \{P_i ::= \bigcup_{j=1}^{m_i} R_{i,j}\}_{i=1}^n$, an unfolding tree of \mathcal{P} rooted at i is a finite tree t such that:

1. each node of t is labeled by a single rule of the system \mathcal{P} ,
2. the root of t is labeled with a rule of P_i ,
3. nodes labeled with base case rules have no successors, and
4. if a node u of t is labeled with a rule whose tail is $P_{i_1}(\mathbf{y}_1) * \dots * P_{i_m}(\mathbf{y}_m)$, then the children of u form the ordered sequence v_1, \dots, v_m where v_j is labeled with one of the rules of P_{i_j} for all $j = 1, \dots, m$.

Remarks. Notice that the recursive predicate $P(x) ::= \exists y. x \mapsto y * P(y)$ does not have finite unfolding trees. However, in general a system of recursive predicates may have infinitely many finite unfolding trees. \square

In the following, we denote by $\mathcal{T}_i(\mathcal{P})$ the set of unfolding trees of \mathcal{P} rooted at i . An unfolding tree $t \in \mathcal{T}_i(\mathcal{P})$ corresponds to a basic formula of separation logic ϕ_t , called the *characteristic formula* of t , and defined in what follows. For a set of tree positions $P \subseteq \mathbb{N}^*$, we denote $LVar^P = \{x^p \mid x \in LVar, p \in P\}$. For a tree position $p \in \mathbb{N}^*$ and a rule R , we denote by R^p the rule obtained by replacing every variable occurrence x in R by x^p . For each position $p \in \text{dom}(t)$, we define a formula ϕ_t^p , by induction on the structure of the subtree of t rooted at p :

- if p is a leaf labeled with a base case rule R , then $\phi_t^p \equiv R^p$
- if p has successors $p.1, \dots, p.m$, and the label of p is the recursive rule $R(\mathbf{x}) \equiv \exists \mathbf{z}. head(R) * \bigstar_{j=1}^m P_{i_j}(\mathbf{y}_j) \wedge pure(R)$, then:

$$\phi_t^p(\mathbf{x}^p) \equiv \exists \mathbf{z}^p. head(R^p) * \bigstar_{j=1}^m [\exists \mathbf{x}_{i_j}^{p.i} \cdot \phi_t^{p.i}(\mathbf{x}_{i_j}^{p.i}) \wedge \mathbf{y}_j^p = \mathbf{x}_{i_j}^{p.i}] \wedge pure(R^p)$$

In the rest of the paper, we write ϕ_t for ϕ_t^ε . Notice that ϕ_t is defined using the set of logical variables $LVar^{dom(t)}$, instead of $LVar$. However the definition of SL semantics from the previous carries over naturally to this case.

Example. (cont'd) Fig. 3(a) presents an unfolding tree for the *tll* predicate given in the previous example. The characteristic formula of each node in the tree can be obtained by composing the formulae labeling the children of the node with the formula labeling the node. The characteristic formula of the tree is the formula of its root. \square

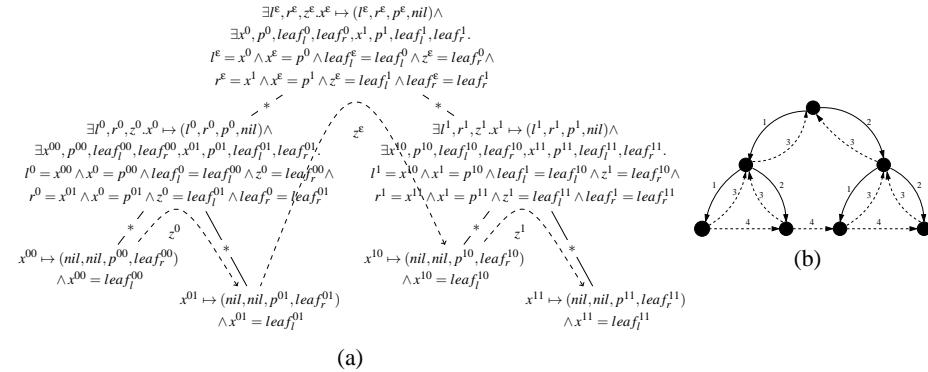


Fig. 3. (a) An unfolding tree for *tll* predicate and (b) a model of the corresponding formula

Given a system of recursive definitions $\mathcal{P} = \{P_i ::= \bigcup_{j=1}^{m_i} R_{i,j}\}_{i=1}^n$, the semantics of a recursive predicate P_i is defined as follows:

$$S, \mathbf{i} \models_{sl} P_i(x_{i,1}, \dots, x_{i,n_i}) \iff S, \mathbf{i}^\varepsilon \models_{sl} \phi_t(x_{i,1}^\varepsilon, \dots, x_{i,n_i}^\varepsilon), \text{ for some } t \in \mathcal{T}_i(\mathcal{P}) \quad (1)$$

where $\mathbf{i}^\varepsilon(x_{i,j}^\varepsilon) \stackrel{\text{def}}{=} \mathbf{i}(x_{i,j})$ for all $j = 1, \dots, n_i$.

Remark. Since the recursive predicate $P(x) ::= \exists y . x \mapsto y * P(y)$ does not have finite unfolding trees, the formula $\exists x.P(x)$ is unsatisfiable. \square

Top Level Formulae. We are now ready to introduce the fragment of *Separation Logic with Recursive Definitions* (SLRD). A formula in this fragment is an existentially quantified formula of the following form: $\exists \mathbf{z} . \varphi * P_{i_1} * \dots * P_{i_n}$, where φ is a basic formula, and P_{i_j} are occurrences of recursive predicates, with free variables in $PVar \cup \mathbf{z}$. The semantics of an SLRD formula is defined in the obvious way, from the semantics of the basic fragment, and that of the recursive predicates.

Example. The following SLRD formulae, with $PVar = \{root, head\}$, describe both the set of binary trees with parent pointer and linked leaves, rooted at *root*, with the leaves

linked into a list pointed to by *head*. The difference is that φ_1 describes also a tree containing only a single allocated location:

$$\begin{aligned}\varphi_1 &\equiv tll(root, nil, head, nil) \\ \varphi_2 &\equiv \exists l, r, x. root \mapsto (l, r, nil, nil) * tll(l, root, head, x) * tll(r, root, x, nil)\end{aligned}\square$$

We are interested in solving two problems on SLD formulae, namely *satisfiability* and *entailment*. The satisfiability problem asks, given a closed SLD formula φ , whether there exists a state S such that $S \models_{sl} \varphi$. The entailment problem asks, given two closed SLD formulae φ_1 and φ_2 , whether for all states S , $S \models_{sl} \varphi_1$ implies $S \models_{sl} \varphi_2$. This is denoted also as $\varphi_1 \models_{sl} \varphi_2$. For instance, in the previous example we have $\varphi_2 \models_{sl} \varphi_1$, but not $\varphi_1 \models_{sl} \varphi_2$.

In general, it is possible to reduce an entailment problem $\varphi_1 \models \varphi_2$ to satisfiability of the formula $\varphi_1 \wedge \neg\varphi_2$. In our case, however, this is not possible directly, because SLD is not closed under negation. The decision procedures for satisfiability and entailment is the subject of the rest of this paper.

3 Decidability of Satisfiability and Entailment in SLD

The decision procedure for the satisfiability and entailment in SLD is based on two ingredients. First, we show that, under certain natural restrictions on the system of recursive predicates, which define a fragment of SLD, called SLD_{btw} , all states that are models of SLD_{btw} formulae have *bounded tree width* (Def. 2). These restrictions are as follows:

1. *Progress*: each rule allocates exactly one variable
2. *Connectivity*: there is at least one selector edge between the variable allocated by a rule and the variable allocated by each of its children in the unfolding tree
3. *Establishment*: all existentially quantified variables in a recursive rule are eventually allocated

Second, we provide a *translation* of SLD_{btw} formulae into equivalent MSO formulae, and rely on the fact that satisfiability of MSO is decidable on classes of states with bounded tree width.

3.1 A Decidable Subset of SLD

At this point we define the SLD_{btw} fragment formally, by defining the three restrictions above. The *progress* condition (1) asks that, for each rule R in the system of recursive definitions, we have $head(R) \equiv \alpha \mapsto (\beta_1, \dots, \beta_n)$, for some variables $\alpha, \beta_1, \dots, \beta_n \in Var_{sl}$. The intuition behind this restriction is reflected by the following example.

Example. Consider the following system of recursive definitions:

$$ls(x, y) ::= x \mapsto y \mid \exists z, t . x \mapsto (z, nil) * t \mapsto (nil, y) * ls(z, t)$$

The predicate $ls(x, y)$ defines the set of structures $\{x(\xrightarrow{1})^n z \mapsto t(\xrightarrow{2})^n y \mid n \geq 0\}$, which clearly cannot be defined in MSO. \square

The *connectivity* condition (2) is defined below:

Definition 4. A rule R of a system of recursive definitions, such that $\text{head}(R) \equiv \alpha \mapsto (\beta_1, \dots, \beta_n)$ and $\text{tail}(R) \equiv \langle P_{i_1}(\mathbf{y}_1), \dots, P_{i_m}(\mathbf{y}_m) \rangle$, $m \geq 1$, is said to be connected if and only if the following hold:

- for each $j = 1, \dots, m$, $(\mathbf{y}_j)_s = \beta'$, for some $1 \leq s \leq n_j$, where n_j is the number of parameters of P_{i_j}
- $\beta_t = \beta'$ occurs in $\text{pure}(R)^*$, for some $1 \leq t \leq n$
- the s -th parameter $x_{i_j,s}$ of P_{i_j} is allocated in the heads of all rules of P_{i_j} .

In this case we say that between rule R and any rule Q of P_{i_j} , there is a *local edge*, labeled by selector t . $\mathcal{F}(R, j, Q) \subseteq \text{Sel}$ denotes the set of all such selectors. If all rules of \mathcal{P} are connected, we say that \mathcal{P} is connected.

Example. The following recursive rule, from the previous tll predicate, is connected:

$$\exists l, r, z . x \mapsto (l, r, p, \text{nil}) * tll(l, x, \text{leaf}_l, z) * tll(r, x, z, \text{leaf}_r) \quad (R_2)$$

R_2 is connected because the variable l is referenced in R_2 and it is passed as the first parameter to tll in the first recursive call to tll . Moreover, the first parameter (x) is allocated by all rules of tll . R_2 is connected, for similar reasons. We have $\mathcal{F}(R_2, 1, R_2) = \{1\}$ and $\mathcal{F}(R_2, 2, R_2) = \{2\}$. \square

The *establishment* condition (3) is formally defined below.

Definition 5. Let $P(x_1, \dots, x_n) = \bigcup_{j=1}^m R_j(x_1, \dots, x_n)$ be a predicate in a recursive system of definitions. We say that a parameter x_i , for some $i = 1, \dots, n$ is allocated in P if and only if, for all $j = 1, \dots, m$:

- either x_i is allocated in $\text{head}(R_j)$, or
- (i) $\text{tail}(R_j) = \langle P_{i_1}(\mathbf{y}_1), \dots, P_{i_k}(\mathbf{y}_k) \rangle$, (ii) $(\mathbf{y}_\ell)_s = x_i$ occurs in $\text{pure}(R_j)^*$, for some $\ell = 1, \dots, k$, and (iii) the s -th parameter of P_{i_ℓ} is allocated in P_{i_ℓ}

A system of recursive definitions is said to be established if and only if every existentially quantified variable is allocated.

Example. Let $\text{llextra}(x) ::= x \mapsto (\text{nil}, \text{nil}) \mid \exists n, e . x \mapsto (n, e) * \text{llextra}(n)$ be a recursive definition system, and let $\phi ::= \text{llextra}(\text{head})$, where $\text{head} \in \text{PVar}$. The models of the formula ϕ are singly-linked lists, where in all locations of the heap, the first selector points to the next location in the list, and the second selector is dangling i.e., it can point to any location in the heap. These dangling selectors may form a squared grid of arbitrary size, which is a model of the formula ϕ . However, the set of squared grids does not have bounded tree width [18]. The problem arises due to the existentially quantified variables e which are never allocated. \square

Given a system \mathcal{P} of recursive definitions, one can effectively check whether it is established, by guessing, for each predicate $P_i(x_{i,1}, \dots, x_{i,n_i})$ of \mathcal{P} , the minimal set of parameters which are allocated in P_i , and verify this guess inductively³. Then, once the minimal set of allocated parameters is determined for each predicate, one can check whether every existentially quantified variable is eventually allocated.

³ For efficiency, a least fixpoint iteration can be used instead of a non-deterministic guess.

Lemma 2. Let $\mathcal{P} = \{P_i ::= \bigcup_{j=1}^{m_i} R_{ij}(x_{i,1}, \dots, x_{i,n_i})\}_{i=1}^k$ be a established system of recursive definitions, and $S = \langle s, h \rangle$ be a state, such that $S, \iota \models_{sl} P_i(x_{i,1}, \dots, x_{i,n_i})$ for some interpretation $\iota : LVar_{sl} \rightarrow_{fin} Loc$ and some $1 \leq i \leq k$. Then $tw(S) \leq \|\mathcal{P}\|^{var}$.

The result of the previous lemma extends to an arbitrary top-level formula:

Theorem 2. Let $\mathcal{P} = \{P_i ::= \bigcup_{j=1}^{m_i} R_{ij}(x_{i,1}, \dots, x_{i,n_i})\}_{i=1}^k$ be a established system of recursive definitions, and $S = \langle s, h \rangle$ be a state, such that $S \models_{sl} \exists \mathbf{z} . \varphi(\mathbf{y}_0) * P_{i_1}(\mathbf{y}_1) * \dots * P_{i_n}(\mathbf{y}_n)$, where φ is a basic SL formula, and P_{i_j} are predicates of \mathcal{P} , and $\mathbf{y}_i \subseteq \mathbf{z}$, for all $i = 0, 1, \dots, n$. Then $tw(S) \leq \max(\|\mathbf{z}\|, |\varphi|, \|PVar\|, \|\mathcal{P}\|^{var})$.

4 From SLRD_{btw} to MSO

This section describes the translation of a SL formula using recursively defined predicates into an MSO formula. We denote by $\Pi(X_0, \dots, X_i, X)$ the fact that X_0, \dots, X_i is a partition of X , and by $\Sigma(x, X)$ the fact that X is a singleton with x as the only element.

4.1 Converting Basic SL Formulae to MSO

For every SL logical variable $x \in LVar_{sl}$ we assume the existence of an MSO logical variable $\bar{x} \in LVar_{mso}$, which is used to replace x in the translation. For every program variable $u \in PVar \setminus \{nil\}$ we assume the existence of a logical variable $\bar{x}_u \in LVar_{mso}$. The special variable $nil \in LVar_{sl}$ is translated into $\bar{x}_{nil} \in LVar_{mso}$ (with the associated MSO constraint $null(\bar{x}_{nil})$). In general, for any pointer or logical variable $\alpha \in Var_{sl}$, we denote by $\bar{\alpha}$, the logical MSO variable corresponding to it.

The translation of a pure SL formula $\alpha = \beta$, $\alpha \neq \beta$, $\pi_1 \wedge \pi_2$ is $\bar{\alpha} = \bar{\beta}$, $\neg(\bar{\alpha} = \bar{\beta})$, $\bar{\pi}_1 \wedge \bar{\pi}_2$, respectively, where $\bar{\pi}(\bar{\alpha}_1, \dots, \bar{\alpha}_k)$ is the translation of $\pi(\alpha_1, \dots, \alpha_k)$. Spatial SL formulae $\sigma(\alpha_1, \dots, \alpha_k)$ are translated into MSO formulae $\bar{\sigma}(\bar{\alpha}_1, \dots, \bar{\alpha}_k, X)$, where X is used for the set of locations allocated in σ . The fact that X actually denotes the domain of the heap, is ensured by the following MSO constraint:

$$Heap(X) \equiv \forall x \bigvee_{i=1}^{\|\text{Sel}\|} (\exists y . edge_i(x, y)) \leftrightarrow X(x)$$

The translation of basic spatial formulae is defined by induction on their structure:

$$\begin{array}{ll} \overline{\text{emp}}(X) & \equiv \forall x . \neg X(x) \\ \overline{(\alpha \mapsto (\beta_1, \dots, \beta_n))}(X) & \equiv \Sigma(\bar{\alpha}, X) \wedge \bigwedge_{i=1}^n edge_i(\bar{\alpha}, \bar{\beta}_i) \wedge \bigwedge_{i=n+1}^{\|\text{Sel}\|} \forall x . \neg edge_i(\bar{\alpha}, x) \\ \overline{(\sigma_1 * \sigma_2)}(X) & \equiv \exists Y \exists Z . \overline{\sigma_1}(Y) \wedge \overline{\sigma_2}(Z) \wedge \Pi(Y, Z, X) \end{array}$$

The translation of a closed basic SL formula φ in MSO is defined as $\exists X . \overline{\varphi}(X)$, where $\overline{\varphi}(X)$ is defined as $\overline{(\pi \wedge \sigma)}(X) \equiv \bar{\pi} \wedge \bar{\sigma}(X)$, and $\overline{(\exists x . \varphi_1)}(X) \equiv \exists \bar{x} . \overline{\varphi_1}(X)$. The following lemma proves that the MSO translation of a basic SL formula defines the same set of models as the original SL formula.

Lemma 3. *For any state $S = \langle s, h \rangle$, any interpretation $\iota : LVar_{sl} \rightarrow_{fin} Loc$, and any basic SL formula φ , we have $S, \iota \models_{sl} \varphi$ if and only if $S, \bar{\iota}, v[X \leftarrow \text{dom}(h)] \models_{mso} \overline{\varphi}(X) \wedge \text{Heap}(X)$, where $\bar{\iota} : LVar_{mso} \rightarrow_{fin} Loc$ is an interpretation of first order variables, such that $\bar{\iota}(x_u) = s(u)$, for all $u \in PVar$, and $\bar{\iota}(\bar{x}) = \iota(x)$, for all $x \in LVar_{sl}$, and $v : LVar_{mso} \rightarrow_{fin} 2^{Loc}$ is any interpretation of second-order variables.*

4.2 States and Backbones

The rest of this section is concerned with the MSO definition of states that are models of recursive SL formulae, i.e. formulae involving recursively defined predicates. The main idea behind this encoding is that any part of a state which is the model of a recursive predicate can be decomposed into a tree-like structure, called the *backbone*, and a set of edges between the nodes in this tree. Intuitively, the backbone is a spanning tree that uses only *local edges*. For instance, in the state depicted in Fig. 3(b), the local edges are drawn in solid lines.

Let $P_k(x_1, \dots, x_n)$ be a recursively defined predicate of a system \mathcal{P} , and $S, \iota \models_{sl} P_k(x_1, \dots, x_n)$, for some state $S = \langle s, h \rangle$ and some interpretation $\iota : LVar_{sl} \rightarrow Loc$. Then $S, \iota \models_{sl} \phi_t$, where $t \in \mathcal{T}_k(\mathcal{P})$ is an unfolding tree, ϕ_t is its characteristic formula, and $\mu : \text{dom}(t) \rightarrow \text{dom}(h)$ is the bijective tree that describes the allocation of nodes in the heap by rules labeling the unfolding tree. Recall that the direction alphabet of the system \mathcal{P} is $\mathcal{D}(\mathcal{P}) = \{-1, 0, \dots, N-1\}$, where N is the maximum number of predicate occurrences within some rule of \mathcal{P} , and denote $\mathcal{D}_+(\mathcal{P}) = \mathcal{D}(\mathcal{P}) \setminus \{-1\}$. For each rule R_{ij} in \mathcal{P} and each direction $d \in \mathcal{D}(\mathcal{P})$, we introduce a second order variable X_{ij}^d to denote the set of locations ℓ such that (i) $t(\mu^{-1}(\ell)) \equiv R_{ij}$ and (ii) $\mu^{-1}(\ell)$ is a d -th child, if $d \geq 0$, or $\mu^{-1}(\ell)$ is the root of t , if $d = -1$. Let $\vec{\mathbf{X}}$ be the sequence of X_{ij}^k variables, enumerated in some order. We use the following shorthands:

$$X_{ij}(x) \equiv \bigvee_{k \in \mathcal{D}(\mathcal{P})} X_{ij}^k(x) \quad X_i(x) \equiv \bigvee_{1 \leq j \leq m_i} X_{ij}(x) \quad X_i^k(x) \equiv \bigvee_{1 \leq j \leq m_i} X_{ij}^k(x)$$

to denote, respectively, locations that are allocated by a rule R_{ij} (X_{ij}), by a recursive predicate P_i (X_i), or by a predicate P_i , who are mapped to a k -th child (or to the root, if $k = -1$) in the unfolding tree of \mathcal{P} , rooted at i (X_i^k).

In order to characterize the backbone of a state, one must first define the local edges:

$$\text{local_edge}_{i,j,p,q}^d(x, y) \equiv \bigwedge_{s \in \mathcal{F}(R_{ij}, d, R_{pq})} \text{edge}_s(x, y)$$

for all $d \in \mathcal{D}_+(\mathcal{P})$. Here $\mathcal{F}(R_{ij}, d, R_{pq})$ is the set of forward local selectors for direction d , which was defined previously – notice that the set of local edges depends on the source and destination rules R_{ij} and R_{pq} , that label the corresponding nodes in the unfolding tree, respectively. The following predicate ensures that these labels are used correctly, and define the successor functions in the unfolding tree:

$$\text{succ}_d(x, y, \vec{\mathbf{X}}) \equiv \bigvee_{\substack{1 \leq i, p \leq M \\ 1 \leq j \leq m_i \\ 1 \leq q \leq m_p}} X_{ij}(x) \wedge X_{pq}^k(y) \wedge \text{local_edge}_{i,j,p,q}^d(x, y)$$

for all $d \in \mathcal{D}_+(\mathcal{P})$. The definition of the backbone of a recursive predicate P_i in MSO follows tightly the definition of the unfolding tree of \mathcal{P} rooted at i (Def. 3):

$$\text{backbone}_i(r, \vec{\mathbf{X}}, T) \equiv \text{tree}(r, \vec{\mathbf{X}}, T) \wedge X_i^{-1}(r) \wedge \text{succ_labels}(\vec{\mathbf{X}})$$

where $\text{tree}(r, \vec{\mathbf{X}}, T)$ defines a tree⁴ with domain T , rooted at r , with successor functions defined by $\text{succ}_0, \dots, \text{succ}_{N-1}$, and succ_labels ensures that the labeling of each tree position (with rules of \mathcal{P}) is consistent with the definition of \mathcal{P} :

$$\text{succ_labels}(\vec{\mathbf{X}}) \equiv \bigwedge_{\substack{1 \leq i \leq M \\ 1 \leq j \leq m_i}} X_{ij}(x) \rightarrow \bigwedge_{d=0}^{r_{ij}-1} \exists y . X_{k_d}^d(y) \wedge \text{succ}_d(x, y, \vec{\mathbf{X}}) \wedge \forall y . \bigwedge_{p=s_{ij}+1}^{\|\text{Sel}\|} \neg \text{edge}_p(x, y)$$

where we suppose that, for each rule R_{ij} of \mathcal{P} , we have $\text{head}(R_{ij}) \equiv \alpha \mapsto (\beta_1, \dots, \beta_{s_{ij}})$ and $\text{tail}(R_{ij}) = \langle P_{k_1}, \dots, P_{k_{r_{ij}}} \rangle$, for some $r_{ij} \geq 0$, and some indexing $k_1, \dots, k_{r_{ij}}$ of predicate occurrences within R_{ij} . The last conjunct ensures that a location allocated in R_{ij} does not have more outgoing edges than specified by $\text{head}(R_{ij})$. This condition is needed, since, unlike SL, the semantics of MSO does not impose strictness conditions on the number of outgoing edges.

4.3 Inner Edges

An edge between two locations is said to be *inner* if both locations are allocated in the heap. Let μ be the bijective tree defined in Sec. 4.2. The existence of an edge $\ell \xrightarrow{k} \ell'$ in S , between two arbitrary locations $\ell, \ell' \in \text{dom}(\mu)$, is the consequence of:

1. a basic points-to formula $\alpha \mapsto (\beta_1, \dots, \beta_k, \dots, \beta_n)$ that occurs in $\mu(\ell)$
2. a basic points-to formula $\gamma \mapsto (\dots)$ that occurs in $\mu(\ell')$
3. a path $\mu(\ell) = p_1, p_2, \dots, p_{m-1}, p_m = \mu(\ell')$ in t , such that the equalities $\beta_k^{p_1} = \delta_2^{p_2} = \dots = \delta_{m-1}^{p_{m-1}} = \gamma^{p_m}$ are all logical consequences of ϕ_t , for some tree positions $p_2, \dots, p_{m-1} \in \text{dom}(t)$ and some variables $\delta_2, \dots, \delta_{m-1} \in LVar_{st}$.

Notice that the above conditions hold only for inner edges. The (corner) case of edges leading to dangling locations is dealt with in Appendix 4.5.

Example. The existence of the edge from tree position 00 to 01 in Fig. 3(b), is a consequence of the following: (1) $x^{00} \mapsto (\text{nil}, \text{nil}, p^{00}, \text{leaf}_r^{00})$, (2) $x^{01} \mapsto (\text{nil}, \text{nil}, p^{01}, \text{leaf}_r^{01})$, and (3) $\text{leaf}_r^{00} = z^0 = \text{leaf}_l^{01} = x^{01}$. The reason for other dashed edges is similar. \square

The main idea here is to encode in MSO the existence of such paths, in the unfolding tree, between the source and the destination of an edge, and use this encoding to define the edges. To this end, we use a special class of tree automata, called *tree-walking automata* (TWA) to recognize paths corresponding to sequences of equalities occurring within characteristic formulae of unfolding trees.

⁴ For space reasons this definition is deferred to Appendix A.

Tree Walking Automata Given a set of tree directions $\mathcal{D} = \{-1, 0, \dots, N\}$ for some $N \geq 0$, a tree-walking automaton⁵, is a tuple $A = (\Sigma, Q, q_i, q_f, \Delta)$ where Σ is a set of tree node labels, Q is a set of states, $q_i, q_f \in Q$ are the initial and final states, and $\Delta : Q \times (\Sigma \cup \{\text{root}\}) \times (\Sigma \cup \{?\}) \rightarrow 2^Q \times (\mathcal{D} \cup \{\varepsilon\})$ is the (non-deterministic) transition function. A configuration of A is a pair $\langle p, q \rangle$, where $p \in \mathcal{D}^*$ is a tree position, and $q \in Q$ is a state. A run of A over a Σ -labeled tree t is a sequence of configurations $\langle p_1, q_1 \rangle, \dots, \langle p_n, q_n \rangle$, with $p_1, \dots, p_n \in \text{dom}(t)$, such that for all $i = 1, \dots, n - 1$, we have $p_{i+1} = p_i.k$, where either:

1. $p_i \neq \varepsilon$ and $(q_{i+1}, k) \in \Delta(q_i, t(p_i), t(p_i.(-1)))$, for $k \in \mathcal{D} \cup \{\varepsilon\}$
2. $p_i = \varepsilon$ and $(q_{i+1}, k) \in \Delta(q_i, \sigma, ?)$, for $\sigma \in \{t(p_i) \cup \text{root}\}$ and $k \in \mathcal{D} \cup \{\varepsilon\}$

The run is said to be *accepting* if $q_1 = q_i$, $p_1 = \varepsilon$ and $q_n = q_f$.

Routing Automata For a system of recursive definitions $\mathcal{P} = \{P_i(x_{i,1}, \dots, x_{i,n_i}) ::= |_{j=1}^{m_i} R_{ij}(x_{i,1}, \dots, x_{i,n_i})\}_{i=1}^k$, we define the TWA $A_{\mathcal{P}} = (\Sigma_{\mathcal{P}}, Q_{\mathcal{P}}, q_i, q_f, \Delta_{\mathcal{P}})$, where $\Sigma_{\mathcal{P}} = \{R_{ij}^k \mid 1 \leq i \leq k, 1 \leq j \leq m_i, k \in \mathcal{D}(\mathcal{P})\}$, $Q_{\mathcal{P}} = \{q_x^{\text{var}} \mid x \in LVar_{\text{sl}}\} \cup \{q_s^{\text{sel}} \mid s \in Sel\} \cup \{q_i, q_f\}$. The transition function $\Delta_{\mathcal{P}}$ is defined as follows:

1. $(q_i, k), (q_s^{\text{sel}}, \varepsilon) \in \Delta(q_i, \sigma, \tau)$ for all $k \in \mathcal{D}_+(\mathcal{P})$, all $s \in Sel$ and all $\sigma \in \Sigma_{\mathcal{P}} \cup \{\text{root}\}$, $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$ i.e., the automaton first moves downwards choosing random directions, while in q_i , then changes to q_s^{sel} for some non-deterministically chosen selector s .
2. $(q_{\beta_s}^{\text{var}}, \varepsilon) \in \Delta(q_s^{\text{sel}}, R_{ij}^k, \tau)$ and $(q_f, \varepsilon) \in \Delta(q_{\alpha}^{\text{var}}, R_{ij}^k, \tau)$ for all $k \in \mathcal{D}(\mathcal{P})$ and $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$ if and only if $\text{head}(R_{ij}) \equiv \alpha \mapsto (\beta_1, \dots, \beta_s, \dots, \beta_m)$, for some $m > 0$ i.e., when in q_s^{sel} , the automaton starts tracking the destination β_s of the selector s through the tree. The automaton enters the final state when the tracked variable α is allocated.
3. for all $k \in \mathcal{D}_+(\mathcal{P})$, all $\ell \in \mathcal{D}(\mathcal{P})$ and all rules $R_{\ell q}$ of $P_{\ell}(x_{\ell,1}, \dots, x_{\ell,n_{\ell}})$, we have $(q_{x_{\ell,j}}^{\text{var}}, k) \in \Delta(q_{y_j}^{\text{var}}, R_{ij}^l, \tau)$, for all $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$, and $(q_{y_j}^{\text{var}}, -1) \in \Delta(q_{x_{\ell,j}}^{\text{var}}, R_{\ell q}^k, R_{ij}^l)$ if and only if $\text{tail}(R_{ij})_k \equiv P_{\ell}(y_1, \dots, y_{n_{\ell}})$ i.e., the automaton moves down along the k -th direction tracking $x_{\ell,j}$ instead of y_j , when the predicate $P_{\ell}(\mathbf{y})$ occurs on the k -th position in R_{ij} . Symmetrically, the automaton can also move up tracking y_j instead of $x_{\ell,j}$, in the same conditions.
4. $(q_{\beta}^{\text{var}}, \varepsilon) \in \Delta(q_{\alpha}^{\text{var}}, R_{ij}^k, \tau)$ for all $k \in \mathcal{D}(\mathcal{P})$ and all $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$ if and only if $\alpha = \beta$ occurs in $\text{pure}(R_{ij})$ i.e., the automaton switches from tracking α to tracking β when the equality between the two variables occurs in R_{ij} , while keeping the same position in the tree.

The following lemma formalizes the correctness of the TWA construction:

Lemma 4. *Given a system of recursive definitions \mathcal{P} , and an unfolding tree $t \in \mathcal{T}_i(\mathcal{P})$ of \mathcal{P} , rooted at i , for any $x, y \in LVar_{\text{sl}}$ and $p, r \in \text{dom}(t)$, we have $\models_{\text{sl}} \phi_t \rightarrow x^p = y^r$ if and only if $A_{\mathcal{P}}$ has a run from $\langle p, q_x^{\text{var}} \rangle$ to $\langle r, q_y^{\text{var}} \rangle$ over t , where ϕ_t is the characteristic formula of t .*

⁵ This notion of tree-walking automaton is a slightly modified but equivalent to the one in [3]. We give the translation of TWA into the original definition in Appendix B.

To the routing automaton $A_{\mathcal{P}}$ corresponds the MSO formula $\Phi_{A_{\mathcal{P}}}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$, where r maps to the root of the unfolding tree, $\vec{\mathbf{X}}$ is the sequence of second order variables X_{ij}^k defined previously, T maps to the domain of the tree, and $\vec{\mathbf{Y}}$ is a sequence of second-order variables X_q , one for each state $q \in Q_{\mathcal{P}}$. We denote by Y_s^{sel} and Y_f the variables from $\vec{\mathbf{Y}}$ that correspond to the states q_s^{sel} and q_f , for all $s \in Sel$, respectively. For space reasons, the definition of $\Phi_{A_{\mathcal{P}}}$ is given in Appendix C. With this notation, we define:

$$inner_edges(r, \vec{\mathbf{X}}, T) \equiv \forall x \forall y \bigwedge_{s \in Sel} \exists \vec{\mathbf{Y}} . \Phi_{A_{\mathcal{P}}}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}}) \wedge Y_s^{sel}(x) \wedge Y_f(y) \rightarrow edge_s(x, y)$$

4.4 Double Allocation

In order to translate the definition of a recursively defined SL predicate $P(x_1, \dots, x_n)$ into an MSO formula \bar{P} , that captures the models of P , we need to introduce a sanity condition, imposing that recursive predicates which establish equalities between variables allocated at different positions in the unfolding tree, are unsatisfiable, due to the semantics of the separating conjunction of SL, which implicitly conjoins all local formulae of an unfolding tree. A double allocation occurs in the unfolding tree t if and only if there exist two distinct positions $p, q \in dom(t)$ and:

1. a basic points-to formula $\alpha \mapsto (\dots)$ occurring in $t(p)$
2. a basic points-to formula $\beta \mapsto (\dots)$ occurring in $t(q)$
3. a path $p = p_1, \dots, p_m = q$ in t , such that the equalities $\alpha^p = \gamma_2^{p_2} = \dots = \gamma_{m-1}^{p_{m-1}} = \beta^q$ are all logical consequences of ϕ_t , for some tree positions $p_2, \dots, p_{m-1} \in dom(t)$ and some variables $\gamma_2, \dots, \gamma_{m-1} \in LVar_{sl}$

The cases of double allocation can be recognized using a routing automaton $B_{\mathcal{P}} = (\Sigma_{\mathcal{P}}, Q'_{\mathcal{P}}, q_i, q_f, \Delta'_{\mathcal{P}})$, whose states $Q'_{\mathcal{P}} = \{q_x^{var} \mid x \in LVar_{sl}\} \cup \{q_0, q_i, q_f\}$ and transitions $\Delta'_{\mathcal{P}}$ differ from $A_{\mathcal{P}}$ only in the following rules:

- $(q_0, \varepsilon) \in \Delta(q_i, \sigma, \tau)$ for all $\sigma \in \Sigma_{\mathcal{P}} \cup \{root\}$ and all $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$, i.e. after nondeterministically choosing a position in the tree, the automaton enters a designated state q_0 , which occurs only once in each run.
- $(q_\alpha^{var}, \varepsilon) \in \Delta(q_0, R_{ij}^k, \tau)$ for all $k \in \mathcal{D}(\mathcal{P})$ and all $\tau \in \Sigma_{\mathcal{P}} \cup \{?\}$ if and only if $head(R_{ij}) = \alpha \mapsto (\dots)$, while in the designated state q_0 , the automaton starts tracking the variable α , which is allocated at that position.

This routing automaton has a run over t , which labels one position by q_0 and a distinct one by q_f if and only if two positions in t allocate the same location. Notice that $B_{\mathcal{P}}$ has always a trivial run that starts and ends in the same position – since each position $p \in dom(t)$ allocates a variable α , and $\langle q_i, \varepsilon \rangle, \dots, \langle q_0, p \rangle, \langle q_\alpha^{var}, p \rangle, \langle q_f, p \rangle$ is a valid run of $B_{\mathcal{P}}$. The predicate system has no double allocation if and only if these are the only possible runs of $B_{\mathcal{P}}$.

The existence of a run of $B_{\mathcal{P}}$ is captured by an MSO formula $\Phi_{B_{\mathcal{P}}}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$, where r maps to the root of the unfolding tree, $\vec{\mathbf{X}}$ is the sequence of second order variables X_{ij}^k defined previously, T maps to the domain of the tree, and $\vec{\mathbf{Y}}$ is the sequence

of second-order variables Y_q , taken in some order, each of which maps to the set of tree positions visited by the automaton while in state $q \in Q'_\mathcal{P}$ – we denote by Y_0 and Y_f the variables from $\vec{\mathbf{Y}}$ that correspond to the states q_0 and q_f , respectively. Finally, we define the constraint: $\text{no_double_alloc}(r, \vec{\mathbf{X}}, T) \equiv \forall \vec{\mathbf{Y}} . \Phi_{B_\mathcal{P}}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}}) \rightarrow Y_0 = Y_f$

4.5 Handling Parameters

The last issue to be dealt with is the role of the actual parameters passed to a recursively defined predicate $P_i(x_{i,1}, \dots, x_{i,k})$ of \mathcal{P} , in a top-level formula. Then, for each parameter $x_{i,j}$ of P_i and each unfolding tree $t \in \mathcal{T}_i(\mathcal{P})$, there exists a path $\varepsilon = p_1, \dots, p_m \in \text{dom}(t)$ and variables $\alpha_1, \dots, \alpha_m \in LVar_{sl}$ such that $x_{i,j} \equiv \alpha_1$ and $\alpha_\ell^{p_\ell} = \alpha_{\ell+1}^{p_{\ell+1}}$ is a consequence of ϕ_t , for all $\ell = 1, \dots, m-1$. Subsequently, there are three (not necessarily disjoint) possibilities:

1. $\text{head}(t(p_m)) \equiv \alpha_m \mapsto (\dots)$, i.e. α_m is allocated
2. $\text{head}(t(p_m)) \equiv \beta \mapsto (\gamma_1, \dots, \gamma_p, \dots, \gamma_\ell)$, and $\alpha_m \equiv \gamma_p$, i.e. α_m is referenced
3. $\alpha_m \equiv x_{i,q}$ and $p_m = \varepsilon$, for some $1 \leq q \leq k$, i.e. α_m is another parameter $x_{i,q}$

Again, we use slightly modified routing automata (one for each of the case above) $C_{\mathcal{P},c}^{i,j} = (\Sigma_\mathcal{P}, Q''_\mathcal{P}, q_i, q_f, \Delta_c^{i,j})$ for the cases $c = 1, 2, 3$, respectively. Here $Q''_\mathcal{P} = \{q_x^{var} \mid x \in LVar_{sl}\} \cup \{q_s^{sel} \mid s \in Sel\} \cup \{q^{i,a} \mid 1 \leq a \leq k\} \cup \{q_i, q_f\}$ and $\Delta_c^{i,j}$, $c = 1, 2, 3$ differ from the transitions of $A_\mathcal{P}$ in the following:

- $(q^{i,j}, \varepsilon) \in \Delta_x^{i,j}(q_i, \text{root}, ?)$, i.e. the automaton marks the root of the tree with a designated state $q^{i,j}$, that occurs only once on each run
- $(q_{x_{i,j}}^{var}, \varepsilon) \in \Delta_x^{i,j}(q^{i,j}, R_{ik}^{-1}, ?)$, for each rule R_{ik} of P_i , i.e. the automaton starts tracking the parameter variable $x_{i,j}$ beginning with the root of the tree
- $(q_f, \varepsilon) \in \Delta_1^{i,j}(q_\alpha^{var}, R_{ij}^k, \tau)$, for all $k \in \mathcal{D}(\mathcal{P})$, $\tau \in \Sigma_\mathcal{P} \cup \{?\}$ iff $\text{head}(R_{ij}) \equiv \alpha \mapsto (\dots)$ is the final rule for $C_{\mathcal{P},1}^{i,j}$
- $(q_s^{sel}, \varepsilon) \in \Delta_2^{i,j}(q_\gamma^{var}, R_{ij}^k, \tau)$, for all $k \in \mathcal{D}(\mathcal{P})$ and $\tau \in \Sigma_\mathcal{P} \cup \{?\}$ iff $\text{head}(R_{ij}) \equiv \alpha \mapsto (\beta_1, \dots, \beta_s, \dots, \beta_n)$ and $\gamma \equiv \beta_s$ i.e., q_s^{sel} is reached in the second case, when the tracked variable is referenced. After that, $C_{\mathcal{P},2}^{i,j}$ moves to the final state i.e., $(q_f, \varepsilon) \in \Delta_2^{i,j}(q_s^{sel}, \sigma, \tau)$ for all $s \in Sel$, all $\sigma \in \Sigma_\mathcal{P} \cup \{\text{root}\}$ and $\tau \in \Sigma_\mathcal{P} \cup \{?\}$
- $(q^{i,a}, \varepsilon) \in \Delta_3^{i,j}(q_{x_{i,a}}^{var}, \text{root}, ?)$ and $(q_f, \varepsilon) \in \Delta_3^{i,j}(q_{i,a}, \text{root}, ?)$, for each $1 \leq a \leq k$ and $a \neq j$ i.e., are the final moves for $C_{\mathcal{P},3}^{i,j}$

The outcome of this construction are MSO formulae $\Phi_{C_{\mathcal{P},c}^{i,j}}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$, for $c = 1, 2, 3$, where r maps to the root of the unfolding tree, respectively, $\vec{\mathbf{X}}$ is the sequence of second order variables X_{ij}^k defined previously, T maps to the domain of the tree, and $\vec{\mathbf{Y}}$ is the sequence of second order variables corresponding to states of $Q''_\mathcal{P}$ – we denote by $Y_f, Y^{i,a}, Y_s^{sel} \in \vec{\mathbf{Y}}$ the variables corresponding to the states $q_f, q^{i,a}$, and q_s^{sel} , respectively.

The parameter $x_{i,j}$ of P_i is assigned by the following MSO constraints:

$$\begin{aligned} param_{i,j}^1(r, \vec{\mathbf{X}}, T) &\equiv \exists \vec{\mathbf{Y}} . \Phi_{C_{\mathcal{P},1}^{i,j}} \wedge Y_0^{i,j}(\bar{x}_{i,j}) \wedge \forall y . Y_f(y) \rightarrow \bar{x}_{i,j} = y \\ param_{i,j}^2(r, \vec{\mathbf{X}}, T) &\equiv \exists \vec{\mathbf{Y}} . \Phi_{C_{\mathcal{P},2}^{i,j}} \wedge Y_0^{i,j}(\bar{x}_{i,j}) \wedge \bigwedge_{s \in Sel} \forall y . Y_s^{sel}(y) \rightarrow edge_s(y, \bar{x}_{i,j}) \\ param_{i,j}^3(r, \vec{\mathbf{X}}, T) &\equiv \exists \vec{\mathbf{Y}} . \Phi_{C_{\mathcal{P},3}^{i,j}} \wedge Y_0^{i,j}(\bar{x}_{i,j}) \wedge \bigwedge_{1 \leq a \leq k} \forall y . Y^{i,a}(y) \rightarrow \bar{x}_{i,j} = \bar{x}_{i,a} \end{aligned}$$

where $\bar{x}_{i,j}$ is the first-order MSO variable corresponding to the SL parameter $x_{i,j}$. Finally, the constraint $param_{i,j}$ is conjunction of the $param_{i,j}^c$, $c = 1, 2, 3$ formulae.

4.6 Translating Top Level $SLRD_{btw}$ Formulae to MSO

We define the MSO formula corresponding to a predicate $P_i(x_{i,1}, \dots, x_{i,n_i})$, of a system of recursive definitions $\mathcal{P} = \{P_1, \dots, P_n\}$:

$$\overline{P}_i(\bar{x}_{i,1}, \dots, \bar{x}_{i,n_i}, T) \equiv \exists r \exists \vec{\mathbf{X}} . backbone_i(r, \vec{\mathbf{X}}, T) \wedge inner_edges(r, \vec{\mathbf{X}}, T) \wedge no_double_alloc(r, \vec{\mathbf{X}}, T) \wedge \bigwedge_{1 \leq j \leq n_i} param_{i,j}(r, \vec{\mathbf{X}}, T)$$

The following lemma is needed to establish the correctness of our construction.

Lemma 5. *For any state $S = \langle s, h \rangle$, any interpretation $\iota : LVar_{sl} \rightarrow_{fin} Loc$, and any recursively defined predicate $P_i(x_1, \dots, x_n)$, we have $S, \iota \models_{sl} P_i(x_1, \dots, x_n)$ if and only if $S, \bar{\iota}, v[T \leftarrow dom(h)] \models_{mso} \overline{P}_i(\bar{x}_1, \dots, \bar{x}_k, T) \wedge Heap(T)$, where $\bar{\iota} : LVar_{mso} \rightarrow_{fin} Loc$ is an interpretation of first order variables, such that $\bar{\iota}(x_u) = s(u)$, for all $u \in PVar$, and $\bar{\iota}(\bar{x}) = \iota(x)$, for all $x \in LVar_{sl}$, and $v : LVar_{mso} \rightarrow_{fin} 2^{Loc}$ is any interpretation of second-order variables.*

Recall that a top level $SLRD_{btw}$ formula is of the form: $\varphi \equiv \exists \mathbf{z} . \phi(\mathbf{y}_0) * P_{i_1}(\mathbf{y}_1) * \dots * P_{i_k}(\mathbf{y}_k)$, where $1 \leq i_1, \dots, i_k \leq n$, and $\mathbf{y}_j \subseteq \mathbf{z}$, for all $j = 0, 1, \dots, k$. We define the MSO formula:

$$\overline{\varphi}(X) \equiv \exists \bar{\mathbf{z}} \exists X_0, \dots, X_k . \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \wedge \overline{P}_{i_1}(\bar{\mathbf{y}}_1, X_1) \wedge \dots \wedge \overline{P}_{i_k}(\bar{\mathbf{y}}_k, X_k) \wedge \Pi(X_0, X_1, \dots, X_k, X)$$

Theorem 3. *For any state S and any closed $SLRD_{btw}$ formula φ we have that $S \models_{sl} \varphi$ if and only if $S \models_{mso} \exists X . \overline{\varphi}(X) \wedge Heap(X)$.*

Theorem 2 and the above theorem prove decidability of satisfiability and entailment problems for $SLRD_{btw}$, by reduction to MSO over states of bounded tree width.

5 Conclusions and Future Work

We defined a fragment of Separation Logic with Recursive Definitions, capable of describing general unbounded mutable data structures, such as trees with parent pointers and linked leaves. The logic is shown to be decidable for satisfiability and entailment, by reduction to MSO over graphs of bounded tree width. We conjecture that the complexity of the decision problems for this logic is elementary, and plan to compute tight upper bounds, in the near future.

Acknowledgement. This work was supported by the Czech Science Foundation (project P103/10/0306) and French National Research Agency (project VERIDYC ANR-09-SEGI-016). We also acknowledge Tomáš Vojnar, Lukáš Holík and the anonymous reviewers for their valuable comments.

References

1. Berdine, J., Calcagno, C., Cook, B., Distefano, D., O’Hearn, P., Wies, T., Yang, H.: Shape analysis for composite data structures. In: Proc. CAV’07. LNCS, vol. 4590. Springer (2007)
2. Berdine, J., Calcagno, C., O’Hearn, P.W.: A decidable fragment of separation logic. In: Proc. of FSTTCS’04. LNCS, vol. 3328. Springer (2004)
3. Bojanczyk, M.: Tree-walking automata. In: Proc. of LATA’08. LNCS, vol. 5196. Springer (2008)
4. Bouajjani, A., Dragoi, C., Enea, C., Sighireanu, M.: A logic-based framework for reasoning about composite data structures. In: Proc. of CONCUR’09. LNCS, vol. 5710. Springer (2009)
5. Bozga, M., Iosif, R., Perarnau, S.: Quantitative separation logic and programs with lists. J. Autom. Reasoning 45(2), 131–160 (2010)
6. Brotherston, J., Kanovich, M.: Undecidability of propositional separation logic and its neighbours. In: Proceedings of the 2010 25th Annual IEEE Symposium on Logic in Computer Science. pp. 130–139. LICS ’10 (2010)
7. Calcagno, C., Distefano, D.: Infer: An automatic program verifier for memory safety of c programs. In: Proc. of NASA Formal Methods’11. LNCS, vol. 6617. Springer (2011)
8. Cook, B., Haase, C., Ouaknine, J., Parkinson, M.J., Worrell, J.: Tractable reasoning in a fragment of separation logic. In: Proc. of CONCUR’11. LNCS, vol. 6901. Springer (2011)
9. Dudka, K., Peringer, P., Vojnar, T.: Predator: A practical tool for checking manipulation of dynamic data structures using separation logic. In: Proc. of CAV’11. LNCS, vol. 6806. Springer (2011)
10. Enea, C., Saveluc, V., Sighireanu, M.: Compositional invariant checking for overlaid and nested linked lists. In: Proc. of ESOP’13. pp. 129–148 (2013)
11. Immerman, N., Rabinovich, A.M., Reps, T.W., Sagiv, S., Yorsh, G.: The boundary between decidability and undecidability for transitive-closure logics. In: Proc of CSL’04. LNCS, vol. 3210. Springer (2004)
12. Iosif, R., Rogalewicz, A., Simacek, J.: The tree width of separation logic with recursive definitions. CoRR abs/1301.5139 (2013)
13. Madhusudan, P., Parlato, G.: The tree width of auxiliary storage. In: Proc. of POPL’11. ACM (2011)
14. Madhusudan, P., Parlato, G., Qiu, X.: Decidable logics combining heap structures and data. In: Proc. of POPL’11 (2011)
15. Møller, A., Schwartzbach, M.I.: The pointer assertion logic engine. In: Proc. of PLDI’01 (June 2001)
16. Nguyen, H.H., Chin, W.N.: Enhancing program verification with lemmas. In: Proc of CAV’08. LNCS, vol. 5123. Springer (2008)
17. Reynolds, J.: Separation Logic: A Logic for Shared Mutable Data Structures. In: Proc. of LICS’02. IEEE CS Press (2002)
18. Seese, D.: The structure of models of decidable monadic theories of graphs. Annals of Pure and Applied Logic 53(2), 169–195 (1991)
19. Yorsh, G., Rabinovich, A.M., Sagiv, M., Meyer, A., Bouajjani, A.: A logic of reachable patterns in linked data-structures. In: Proc. of FoSSaCS’06. LNCS, vol. 3921. Springer (2006)

A Definition of tree structures in MSO

Let $\vec{\mathbf{X}} = \{X_1, \dots, X_m\}$ define a set of tree labels. Given a direction alphabet $\mathcal{D} = \{-1, 0, \dots, N\}$, we consider a set of (partial) successor functions $Succ_{\mathcal{D}} = \{succ_0, \dots, succ_N\}$. These functions can be encoded by MSO formulae $succ_d(x, y, \vec{\mathbf{X}})$, and are supposed to satisfy the following constraint:

$$\forall x, y, z . \bigwedge_{i \in \mathcal{D}_+} succ_i(x, y) \wedge succ_i(x, z) \rightarrow y = z$$

A tree structure with root x , domain X and successor functions $Succ_{\mathcal{D}}$ is defined by the MSO formula $tree(r, \vec{\mathbf{X}}, T)$, which is the conjunction of $\Pi(\vec{\mathbf{X}}, T)$ and the following four MSO constraints:

(A) x is the root of the tree:

$$\forall y . X(y) \rightarrow \bigwedge_{d \in \mathcal{D}_+} \neg succ_d(y, x, \vec{\mathbf{X}})$$

(B) successors are pairwise distinct:

$$\forall x, y, z . \bigwedge_{0 \leq i < j \leq N} succ_i(x, y, \vec{\mathbf{X}}) \wedge succ_j(x, z, \vec{\mathbf{X}}) \rightarrow y \neq z$$

(C) each node except for the root has exactly one predecessor:

$$\forall y . x \neq y \rightarrow \exists! z . \bigvee_{i \in \mathcal{D}_+} succ_i(z, y)$$

where $\exists!$ stands for the unique existential quantification

(D) all nodes in X are reachable from x :

$$\begin{aligned} closed(X) &\equiv \forall y . X(y) \wedge \bigwedge_{i \in \mathcal{D}_+} \exists z . succ_i(y, z) \rightarrow X(z) \\ reach(x, X) &\equiv X(x) \wedge closed(X) \wedge \forall Y . Y(x) \wedge closed(Y) \rightarrow X \subseteq Y \end{aligned}$$

The following lemma formalizes the correctness of this definition.

Lemma 6. *For any state $S = \langle s, h \rangle$ and interpretations $\iota : LVar_{mso} \rightarrow_{fin} Loc$ and $\nu : LVar_{mso} \rightarrow_{fin} 2^{Loc}$ such that $\iota(r) \in \text{dom}(h)$ and $\nu(X_i) \subseteq \text{dom}(h)$, for all $i = 1, \dots, m$, and $\nu(T) = \text{dom}(h)$, we have:*

$$S, \iota, \nu \models_{mso} tree(r, \vec{\mathbf{X}}, T)$$

if and only if there exists a unique prefix-closed set $P \subseteq \mathbb{N}^$ and two unique trees $\mu : P \rightarrow \text{dom}(h)$ and $\lambda : P \rightarrow \{X_1, \dots, X_m\}$, such that μ is bijective, $\mu(\varepsilon) = \iota(r)$, and:*

1. $\forall \ell \in \text{dom}(h) \forall d \in \mathcal{D}_+ . S, \iota[x \leftarrow \ell][y \leftarrow \mu(\ell)], \nu \models_{mso} \neg succ_d(x, y, \vec{\mathbf{X}})$
2. $\forall p \in P \forall d \in \mathcal{D}_+ . S, \iota[x \leftarrow \mu(p)][y \leftarrow \mu(p.d)], \nu \models_{mso} succ_d(x, y, \vec{\mathbf{X}}) \iff p.d \in D$
3. $\nu(X_i) = \{\ell \in \text{dom}(h) \mid \lambda(\mu^{-1}(\ell)) = X_i\}$

Proof. “ \Rightarrow ” Since $S, \iota, v \models_{mso} \Pi(\vec{\mathbf{X}}, T)$, we have that $v(X_1), \dots, v(X_m)$ forms a partition of $dom(h)$. We define D , μ and λ , as the limits of the increasing sequences defined as follows. Let $P_0 = \{\varepsilon\}$, $\mu_0 = \{(\varepsilon, \iota(r))\}$ and $\lambda_0 = \{(\varepsilon, X_j)\}$, where $j \in \{1, \dots, m\}$ is the unique index such that $\iota(r) \in v(X_j)$, and, for all $i \geq 0$:

- $P_{i+1} = P_i \cup \{p.d \mid p \in P_i, d \in \mathcal{D}_+, p.d \notin P_i, \exists \ell \in dom(h) . S, \iota[x \leftarrow \mu_i(p)][y \leftarrow \ell], v \models_{mso} succ_d(x, y, \vec{\mathbf{X}})\}$
- $\mu_{i+1} = \mu_i \cup \{(p.d, \ell) \mid p \in D_i, d \in \mathcal{D}_+, p.d \notin P_i, S, \iota[x \leftarrow \mu_i(p)][y \leftarrow \ell], v \models_{mso} succ_d(x, y, \vec{\mathbf{X}})\}$; notice that the choice of ℓ is unique, because $succ_d$ defines a partial function
- $\lambda_{i+1} = \lambda_i \cup \{(p, X_j) \mid p \in dom(\mu_{i+1}) \setminus dom(\mu_i), \mu_{i+1}(p) \in v(X_j)\}$, where $j \in \{1, \dots, m\}$ is the unique index such that $\mu_{i+1}(p) \in v(X_j)$

The sequences stabilize, because $dom(h)$ is finite, and we define $P = \bigcup_{i \geq 0} P_i$, $\mu = \bigcup_{i \geq 0} \mu_i$ and $\lambda = \bigcup_{i \geq 0} \lambda_i$. Moreover $\mu_0(\varepsilon) = \iota(r)$, hence $\mu(\varepsilon) = \iota(r)$. The first condition is satisfied as a consequence of point (A) in the definition of $tree(r, \vec{\mathbf{X}}, T)$, and the second condition can be proved by induction on the definitions of P_i and μ_i . The fact that μ is bijective is a consequence of points (B), (C) and (D) in the definition of $tree(r, \vec{\mathbf{X}}, T)$. First, suppose that μ is not one-to-one, i.e. there exist two distinct positions $p, q \in dom(\mu)$ such that $\mu(p) = \mu(q) = \ell$. Since $p \neq q$, either:

- (i) p is a prefix of q , or viceversa
- (ii) there exist a position r such that $r.d_1$ is a prefix of p and $r.d_2$ is a prefix of q , for some $d_1, d_2 \in \mathcal{D}_+$, $d_1 \neq d_2$.

In both cases we obtain a contradiction. Second, suppose that μ is not onto, i.e. there exists $\ell \in dom(h)$ such that $\mu(p) \neq \ell$, for all $p \in dom(\mu)$. But this is clearly in contradiction with point (D) above and the definition of D and μ . The third condition can be proved inductively on the definition of λ_i . Finally P , μ and λ are unique, since the choices at each step $i \geq 0$ in the definition of P_i , μ_i and λ_i are unique. “ \Leftarrow ” This direction is an easy exercise. \square

B Standard Tree Walking Automata

We recall the standard definition of a TWA from [3]. Given a set of tree directions $\mathcal{D} = \{-1, 0, \dots, N\}$, for some $N \geq 0$, a *standard tree-walking automaton* (STWA) is a tuple $A_s = (\Sigma, Q, q_i, q_f, \Delta_s)$ where Σ is a set of tree node labels, Q is a set of states, $q_i, q_f \in Q$ are the initial and final states, and $\Delta_s : Q \times (\mathcal{D}_+ \cup \{root\}) \times \Sigma \rightarrow 2^Q \times (\mathcal{D} \cup \{\varepsilon\})$ is the (non-deterministic) transition function. A configuration of A is a pair $\langle p, q \rangle$, where $p \in \mathcal{D}_+^*$ is a tree position, and $q \in Q$ is a state. A run of A over a Σ -labeled tree t is a sequence of configurations $\langle p_1, q_1 \rangle, \dots, \langle p_n, q_n \rangle$, with $p_1, \dots, p_n \in dom(t)$ and $q_1, \dots, q_n \in Q$, such that, for all $i = 1, \dots, n - 1$, we have $p_{i+1} = p_i.k$ for some $k \in \mathcal{D} \cup \{\varepsilon\}$, where either:

1. $p_i = p.d$ for some $p \in \mathcal{D}_+^*$, $d \in \mathcal{D}_+$ and $(q_{i+1}, k) \in \Delta(q_i, d, t(p_i))$
2. $p_i = \varepsilon$ and $(q_{i+1}, k) \in \Delta(q_i, root, t(p_i))$

The run is said to be *accepting* if $q_1 = q_i$, $p_1 = \varepsilon$ and $q_n = q_f$, in which case we say that A accepts t . We denote by $\mathcal{L}(A)$ the set of trees accepted by a (S)TWA A .

Lemma 7. *For each TWA $A = (\Sigma, Q, q_i, q_f, \Delta)$ there exists an STWA $A_s = (\Sigma, Q \cup Q_{aux}, q_i, q_f, \Delta_s)$ such that $\mathcal{L}(A) = \mathcal{L}(A_s)$.*

Proof. For each rule in Δ we create a set of rules in Δ_s , such that $\langle p_1, q_1 \rangle \xrightarrow{A} \langle p_2, q_2 \rangle$ if and only if $\langle p_1, q_1 \rangle \xrightarrow{\Delta_s^+} \langle p_2, q_2 \rangle$, i.e. we simulate the effect of a single step in A by a sequence of steps in A_s . The construction of Δ_s is done as follows. Let $(q_j, k) \in \Delta(q_i, \sigma, \pi)$ be a transition rule of A .

- if $\sigma = root$ and $\pi = ?$ then, for each $\tau \in \Sigma$, we have $(q_j, k) \in \Delta_s(q_i, root, \tau)$
- if $\sigma \in \Sigma$ and $\pi = ?$ then $(q_j, k) \in \Delta_s(q_i, root, \sigma)$
- if $\sigma, \pi \in \Sigma$ then, for each $d \in \mathcal{D}_+$, we have the following sequence of rules:
 - $(q_{i,d}^1, -1) \in \Delta_s(q_i, d, \sigma)$
 - $(q_{i,d}^2, d) \in \Delta_s(q_{i,d}^1, e, \pi)$, for each $e \in \mathcal{D}_+ \cup \{root\}$
 - $(q_j, k) \in \Delta_s(q_{i,d}^2, d, \sigma)$
 for two fresh states $q_{i,m}^1, q_{i,m}^2 \in Q_{aux} \setminus Q$.

The following proofs are left as an easy exercise.

1. each sequence of steps $\langle p_1, q_1 \rangle, \dots, \langle p_n, q_n \rangle$ of A corresponds to a unique sequence of steps of A_s starting and ending in the same configurations
2. each sequence of steps $\langle p_1, q_1 \rangle, \dots, \langle p_n, q_n \rangle$ of A_s corresponds to a unique sequence of steps of A starting and ending in the same configurations

□

C MSO encoding of Tree Walking Automata

We consider a class of tree structures with successor functions $Succ_{\mathcal{D}} = \{succ_i \mid i \in \mathcal{D}_+\}$ for some set of directions $\mathcal{D} = \{-1, 0, \dots, N\}$, $N \geq 0$, with labels from the alphabet $\vec{\mathbf{X}} = \langle X_1, \dots, X_m \rangle$ of second-order variables. These labels define a partition on the domain of the tree, i.e. we assume that the following constraint holds in what follows:

$$tree(r, \vec{\mathbf{X}}, T)$$

We define the predecessor function $succ_{-1}$ as follows:

$$succ_{-1}(x, y, \vec{\mathbf{X}}) \equiv \bigvee_{0 \leq i \leq k} succ_i(y, x, \vec{\mathbf{X}})$$

Since the successors of a node in a tree are pairwise distinct, $succ_{-1}$ is a well-defined partial function.

Let $A = (\vec{\mathbf{X}}, Q, q_i, q_f, \Delta)$ be a tree walking automaton, and let $Q = \{q_1, \dots, q_k\}$ be some arbitrary indexing of the set of states. W.l.o.g. we assume that no transition rule

in Δ originates in q_f . Let $\vec{\mathbf{Y}} = \langle Y_1, \dots, Y_k \rangle$ be a sequence of second-order variables, one for each state. First, we define a step $\langle x, q_i \rangle \rightarrow \langle y, q_j \rangle$ of A on the tree, as follows:

$$\text{step}(x, y, \vec{\mathbf{X}}) \equiv \bigvee_{(q_j, k) \in \Delta(q_i, x_p, x_q)} (Y_i(x) \wedge Y_j(y) \wedge X_p(x) \wedge \text{succ}_k(x, y, \vec{\mathbf{X}}) \wedge (x = r \vee \exists z . \text{succ}_{-1}(x, z, \vec{\mathbf{X}}) \wedge X_q(z)))$$

Any position on the run is reachable from the root r , with respect to the *step* relation:

$$\begin{aligned} \text{step_closed}(X) &\equiv \forall x, y . X(x) \wedge \text{step}(x, y, \vec{\mathbf{X}}) \rightarrow X(y) \\ \text{step_reach}(x, X) &\equiv X(x) \wedge \text{step_closed}(X) \wedge \forall Y . Y(x) \wedge \text{step_closed}(Y) \rightarrow X \subseteq Y \end{aligned}$$

The run R of A is defined by the conjunction of the following constraints:

- (A) R equals the union of Y_1, \dots, Y_k i.e., each position in the run is marked by at least one state of the automaton
- (B) the root r of the tree is labeled with Y_i , where Y_i is the second-order variable corresponding to the initial state q_i :

$$Y_i(r)$$

- (C) the final position of x_f the run is labeled with Y_f , where Y_f is the second-order variable corresponding to the final state q_f :

$$\forall y . \neg \text{step}(x_f, y, \vec{\mathbf{X}}) \wedge Y_f(x_f)$$

- (D) every non-final position has a successor position in the run:

$$R \subseteq T \wedge \text{step_reach}(r, R) \wedge \forall x . x \neq x_f \rightarrow \exists y . \text{step}(x, y, \vec{\mathbf{X}})$$

The final formula $\Phi_A(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$ is obtained by conjoining the above constraints and existentially quantifying x_f and R . The following lemma formalizes the correctness of this construction:

Lemma 8. *For any state $S = \langle s, h \rangle$ and interpretations $\iota : L\text{Var}_{\text{mso}} \rightarrow_{\text{fin}} \text{Loc}$ and $\nu : L\text{Var}_{\text{mso}} \rightarrow_{\text{fin}} 2^{\text{Loc}}$ where $\iota(r) \in \text{dom}(h)$ and $\nu(X_i) \subseteq \text{dom}(h)$, for all $i = 1, \dots, m$ and $\nu(T) = \text{dom}(h)$, such that:*

$$S, \iota, \nu \models_{\text{mso}} \text{tree}(r, \vec{\mathbf{X}}, T)$$

let $P \subseteq \mathbb{N}^$ and $\mu : P \rightarrow \text{dom}(h)$, $\lambda : P \rightarrow \{X_1, \dots, X_m\}$ be the prefix-closed set and unique trees from Lemma 6. Then we have:*

$$S, \iota, \nu \models_{\text{mso}} \Phi_A(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$$

if and only if A has a loop-free accepting run π over λ such that $\nu(Y_j) = \{\mu(p) \mid p \in P \text{ and } \langle p, q_j \rangle \text{ occurs on } \pi\}$, for all $j = 1, \dots, k$.

Proof. “ \Rightarrow ” From the definition of Φ_A , we can construct a loop-free maximal path $\ell_0, \ell_1, \dots, \ell_n$ in $\text{dom}(h)$ such that:

- $\iota(r) = \ell_0 \in \nu(Y_i)$ and $\ell_n \in \nu(Y_f)$
- $S, \iota, \nu[x \leftarrow \ell_i][y \leftarrow \ell_{i+1}] \models_{\text{mso}} \text{step}(x, y, \vec{\mathbf{X}})$, for all $i \geq 0$
- $\bigcup_{j=1}^k \nu(Y_j) = \{\ell_0, \ell_1, \dots, \ell_n\}$

Then $\varepsilon = \mu^{-1}(\ell_0), \mu^{-1}(\ell_1), \dots, \mu^{-1}(\ell_n)$ is a path in P and A has a loop-free accepting run $\pi : \langle \mu^{-1}(\ell_0), q_i \rangle, \dots, \langle \mu^{-1}(\ell_n), q_f \rangle$ over λ . Moreover $\langle \mu^{-1}(\ell_i), q_j \rangle$ occurs on the run if and only if $\ell_i \in \nu(Y_j)$. “ \Leftarrow ” This direction is left as an easy exercise. \square

D Routing Automaton Example

Routing automaton for the tree with link leaves—predicate tll from Section 2.2. There is only a single predicate tll with two rules. The routing automaton is $A_{tll} = (\Sigma_{\mathcal{P}}, Q, q_i, q_f, \Delta)$, where

- $\Sigma_{\mathcal{P}} = \{R_{11}^0, R_{12}^0, R_{11}^1, R_{12}^1, R_{11}^{-1}, R_{12}^{-1}\}$
- $Q = \{q_x^{var}, q_p^{var}, q_{leaf_l}^{var}, q_{leaf_r}^{var}, q_l^{var}, q_r^{var}, q_z^{var}\} \cup \{q_1^{sel}, q_2^{sel}, q_3^{sel}, q_4^{sel}\} \cup \{q_i, q_f\}$

and Δ is defined as follows (the numbers corresponds to numbers in definition of routing automata):

1. - $(q_i, k), (q_1^{sel}, \varepsilon), (q_2^{sel}, \varepsilon), (q_3^{sel}, \varepsilon), (q_4^{sel}, \varepsilon) \in \Delta(q_i, \sigma, \tau), \sigma \in \Sigma_{\mathcal{P}}, k \in \{0, 1\}, \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
2. - $(q_l^{var}, \varepsilon) \in \Delta_{init}(q_1^{sel}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_r^{var}, \varepsilon) \in \Delta_{init}(q_2^{sel}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_p^{var}, \varepsilon) \in \Delta_{init}(q_3^{sel}, \sigma, \tau), \sigma \in \Sigma_{\mathcal{P}}, \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_{leaf_r}^{var}, \varepsilon) \in \Delta_{init}(q_4^{sel}, R_{11}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_f, \varepsilon) \in \Delta(q_x^{var}, \sigma, \tau), \sigma \in \Sigma_{\mathcal{P}}, \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
3. - $(q_x^{var}, 0) \in \Delta(q_l^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_l^{var}, -1) \in \Delta(q_x^{var}, \sigma^0, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_x^{var}, 1) \in \Delta(q_r^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_r^{var}, -1) \in \Delta(q_x^{var}, \sigma^1, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_p^{var}, 0), (q_p^{var}, 1) \in \Delta(q_x^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_x^{var}, -1) \in \Delta(q_p^{var}, \sigma, \tau), \sigma \in \Sigma_{\mathcal{P}}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_{leaf_r}^{var}, 0), (q_{leaf_r}^{var}, 1) \in \Delta(q_z^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_z^{var}, -1) \in \Delta(q_{leaf_r}^{var}, \sigma^1, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_z^{var}, -1) \in \Delta(q_{leaf_r}^{var}, \sigma^0, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_{leaf_l}^{var}, 0) \in \Delta(q_{leaf_l}^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_{leaf_l}^{var}, -1) \in \Delta(q_{leaf_l}^{var}, \sigma^0, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
- $(q_{leaf_l}^{var}, 1) \in \Delta(q_{leaf_l}^{var}, R_{12}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_{leaf_l}^{var}, -1) \in \Delta(q_{leaf_l}^{var}, \sigma^1, \tau), \sigma \in \{R_{11}, R_{12}\}, \tau \in \Sigma_{\mathcal{P}}$
4. - $(q_x^{var}, \varepsilon) \in \Delta(q_{leaf_l}^{var}, R_{11}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$
- $(q_{leaf_l}^{var}, \varepsilon) \in \Delta(q_x^{var}, R_{11}^k, \tau), k \in \mathcal{D}(\mathcal{P}), \tau \in \Sigma_{\mathcal{P}} \cup \{?\}$

E Missing Proofs

Proof of Theorem 1: Each state $S = \langle s, h \rangle$ is both a vertex- and edge-labeled graph, in whose set of vertices is $loc(S)$, vertices are labeled with pointer variables, and edges with selectors. There are three MSO-definable restrictions making the difference between states and arbitrary graphs:

1. each pointer variable $u \in PVar$ labels at most one location:

$$\rho_1 \equiv \forall x.y . \bigwedge_{u \in PVar} var_u(x) \wedge var_u(y) \rightarrow x = y$$

2. each edge leads to at most one location:

$$\rho_2 \equiv \forall x, y, z . \bigwedge_{s \in Sel} edge_s(x, y) \wedge edge_s(x, z) \rightarrow y = z$$

3. there exists a unique designated nil location with no outgoing edges:

$$\rho_3 \equiv \exists !x \forall y . \bigwedge_{s \in Sel} \neg edge_s(x, y) \wedge \forall z . null(y) \rightarrow x = z$$

Hence the satisfiability problem for an MSO formula ϕ interpreted over states is equivalent to the satisfiability of the MSO formula $\phi \wedge \rho_1 \wedge \rho_2 \wedge \rho_3$ interpreted over arbitrary graphs. The latter problem is decidable, as shown, e.g. by Theorem 2.1 in [13]. \square

Proof of Lemma 1: Let $\phi \equiv \Sigma \wedge \Pi$, where Σ and Π are the spatial and pure parts of ϕ , respectively. First observe that $dom(h) = \{(s \oplus \iota)(\alpha) \mid \alpha \mapsto (\dots, \beta, \dots) \text{ occurs in } \Sigma\}$ and $Img(h) = \{(s \oplus \iota)(\beta) \mid \alpha \mapsto (\dots, \beta, \dots) \text{ occurs in } \Sigma\}$. Hence $|\phi| = |\Sigma| \geq \|dom(h) \cup Img(h)\|$. Moreover, we have that $\|img(s)\| \leq \|PVar\|$.

We define a tree decomposition of S as follows. Let $P \subseteq \mathbb{N}^*$ be a prefix-closed set such that $\|P\| = \|dom(h) \cup Img(h)\|$ and $\delta : \{0.p \mid p \in P\} \cup \{\varepsilon\} \rightarrow 2^{loc(S)}$ be a tree such that $\delta(0.p) = dom(h) \cup Img(h)$, for all $p \in P$ and $\delta(\varepsilon) = img(s) \setminus (dom(h) \cup Img(h))$. It is easy to check that δ satisfies the conditions of Def. 2. Also, $\|\delta(0.p)\| \leq |\phi|$, for all $p \in P$, and $\|\delta(\varepsilon)\| \leq \|PVar\|$ i.e., $\|\delta(p)\| \leq \max(|\phi|, \|PVar\|)$, for all $p \in dom(\delta)$. \square

Proof of Lemma 2: By the definition of the semantics of recursive predicates, we have:

$$S, \iota^\varepsilon \models_{sl} \phi_t(x_{i,1}^\varepsilon, \dots, x_{i,n_i}^\varepsilon)$$

for some unfolding tree $t \in \mathcal{T}_i(\mathcal{P})$. Observe that the only free variables of ϕ_t are $x_{i,1}^\varepsilon, \dots, x_{i,n_i}^\varepsilon$, the rest occurring under existential quantification. Let $\overline{\phi_t}$ be the (matrix) formula obtained from ϕ_t by renaming each existentially quantified variable to a unique name, and forgetting the existential quantifiers. Also let $\overline{\iota^\varepsilon} : LVar^{dom(t)} \rightharpoonup_{fin} Loc$ be an interpretation such that:

$$S, \overline{\iota^\varepsilon} \models_{sl} \overline{\phi_t}$$

By the definition of the semantics of SL, such an interpretation must exist. Hence for each position $p \in dom(t)$, there exists a state $S_p = \langle s, h_p \rangle$ such that:

$$S_p, \overline{\iota^\varepsilon} \models_{sl} head(t(p))$$

and, moreover $\|dom(h_p)\| = 1$, since, by convention, $head(t(p))$ allocates exactly one variable. Consequently, there exists a bijective tree $\mu : dom(t) \rightarrow dom(h)$ such that, for all $p \in dom(t)$, we have $dom(h_p) = \{\mu(p)\}$.

We define a tree decomposition $\delta : dom(t) \rightarrow 2^{Loc}$ as follows, for all positions $p_0 \in dom(t)$, $\delta(p_0)$ contains only the following locations:

- (i) $\mu(p_0) \in \delta(p_0)$
- (ii) $\iota(x_{i,j}) \in \delta(p_0)$, for all $1 \leq j \leq n_i$

- (iii) if $head(t(p_0)) \equiv \alpha \mapsto (\beta_1, \dots, \beta_s)$, for each $u = 1, \dots, s$ and each sequence $\beta_u^{p_0} = \dots = \gamma^{p_1}$ of equalities occurring in ϕ_t , such that $head(t(p_1))$ allocates γ , we have $\mu(p_1) \in \delta(p)$, for each position p within the sequence

First, we prove that δ is a valid tree decomposition (Def. 2):

1. Let $\ell \in loc(S)$ be a location. If $\ell \in dom(h)$, then $\ell \in \delta(\mu^{-1}(\ell))$ by point (i) above. If $\ell \in Img(h) \setminus dom(h)$, since \mathcal{P} is established, then $\ell = \iota(x_{i,j})$, for some $1 \leq j \leq n_i$, by point (ii) above. Hence $\ell \in \delta(p)$, for all $p \in dom(\delta)$. Consequently $loc(S) \subseteq \bigcup_{p \in dom(\delta)} \delta(p)$, and the other direction is trivial.
2. Let $\ell_1 \xrightarrow{s} \ell_2$ be an edge in S . Then $\ell_1 \in dom(h)$ and $\ell_1 \in \delta(\mu^{-1}(\ell_1))$, by point (i) above. But $\ell_1 \xrightarrow{s} \ell_2$ only exists because $head(t(\mu^{-1}(\ell_1))) \equiv \alpha \mapsto (\beta_1, \dots, \beta_s)$, and there exists a sequence $\beta_u^{\mu^{-1}(\ell_1)} = \dots = \gamma^{\mu^{-1}(\ell_2)}$ of equalities occurring in ϕ_t , for some $u = 1, \dots, s$, such that $head(t(\mu^{-1}(\ell_2)))$ allocates γ . By point (iii) above, we have $\ell_2 \in \delta(\mu^{-1}(\ell_1))$.
3. Let $p, r \in dom(\delta)$ be two distinct locations, q be on the path from p to r , and let $\ell \in \delta(p) \cap \delta(r)$. Then there are two possibilities. Either $\ell = \iota(x_{i,j})$, for some $1 \leq j \leq n_i$, in which case $\ell \in \delta(q)$, by point (ii) above. Otherwise, the only remaining possibility is that both p and r are on a sequence of equalities $\alpha_1^{p_1} = \dots = \alpha_m^{p_m}$, and $\ell \in \bigcap_{i=1}^m \delta(p_i)$, cf. point (iii) above. But in this case q must be one the same sequence, hence $\ell \in \delta(q)$.

Finally, we prove that $\|\delta(p)\| \leq \|\mathcal{P}\|^{var}$, for all $p \in dom(\delta)$. Let $\ell \in \delta(p)$ be a location. There are two reasons for $\ell \in \delta(p)$:

- $\ell = \underline{\iota}(x_{i,j})$, for some $j = 1, \dots, n_i$, cf. point (ii) above
- $\ell = \underline{\iota}^e(\alpha^p)$, where α is an existentially quantified variable that occurs within $t(p)$, cf. points (i) and (iii) above

Hence $\|\delta(p)\|$ may not exceed the maximum number of variables that occur either free, or existentially quantified, within $t(p)$. \square

Proof of Theorem 2: Let $\iota : LVar_{sl} \rightarrow_{fin} Loc$ be an interpretation, and $S_0 = \langle s_0, h_0 \rangle, S_1 = \langle s_1, h_1 \rangle, \dots, S_n = \langle s_n, h_n \rangle$ be states such that:

$$\begin{aligned} S_0, \iota &\models_{sl} \varphi \\ S_1, \iota &\models_{sl} P_{i_1} \\ &\dots \\ S_n, \iota &\models_{sl} P_{i_n} \end{aligned}$$

and $S = S_0 \uplus S_1 \uplus \dots \uplus S_n$. By Lemma 1, $tw(S_0) \leq \max(|\varphi|, \|PVar\|)$, and by Lemma 2, $tw(S_i) \leq \|\mathcal{P}\|^{var}$, for all $i = 1, \dots, n$. Hence there exist prefix-closed sets P_i and tree decompositions $\delta_i : P_i \rightarrow loc(S_i)$ of S_i , for all $i = 0, 1, \dots, n$, respectively.

We define the prefix-closed set $P = \{i.p \mid i = 0, 1, \dots, n, p \in P_i\} \cup \{\varepsilon\}$ and a tree decomposition δ of S as follows. Let $\delta(\varepsilon) = \{\iota(x) \mid x \in \mathbf{z}\}$ and $\delta(i.p) = \delta_i(p)$, for all $p \in P_i$, and $i = 0, 1, \dots, n$. Let us first check that δ meets the conditions of Def. 2. The first point follows from the fact that $loc(S) = \bigcup_{i=0}^n loc(S_i)$, and $loc(S_i) = \bigcup_{p \in P_i} \delta_i(p)$,

for all $i = 0, 1, \dots, n$. Second, let $\ell \xrightarrow{s} m$ be an edge in S . Then $\ell \in \text{dom}(h_i)$, for some $i = 0, 1, \dots, n$, and therefore $m \in \text{Img}(h_i)$. Hence there exists $p \in P_i$ such that $\ell, m \in \delta_i(p)$. Third, let $q \in P$ be on a path from p to r , with $p, r \in P$. We distinguish two cases:

- $q = \varepsilon$, $p = i.p'$ and $r = j.r'$, with $p' \in P_i$ and $r' \in P_j$, $0 \leq i < j \leq n$. Then $\delta(p) \cap \delta(r) \subseteq \{\iota(x) \mid x \in \mathbf{z}\} = \delta(q)$
- $q = i.q'$, $q' \in P_i$, and $p = i.p'$, $r = i.r'$ for some $i = 0, 1, \dots, n$. Then $\delta_i(p') \cap \delta_i(r') \subseteq \delta_i(q')$ by the fact that δ_i is a tree decomposition of S_i , and hence $\delta(p) \cap \delta(r) \subseteq \delta(q)$

Finally, $\|\delta(0.p)\| \leq \max(|\varphi|, \|PVar\|)$ for all $p \in P_0$, $\|\delta(i.p)\| \leq \|P\|^{\text{var}}$ for all $p \in P_i$, $i = 1, \dots, n$, and $\|\delta(\varepsilon)\| \leq \|\mathbf{z}\|$. Hence $\|\delta(p)\| \leq \max(\|\mathbf{z}\|, |\varphi|, \|PVar\|, \|P\|^{\text{var}})$, for all $p \in \text{dom}(\delta)$. \square

Proposition 1. *For any state $S = \langle s, h \rangle$ and any interpretations $\iota : LVar_{mso} \rightarrow_{fin} Loc$ and $v : LVar_{mso} \rightarrow_{fin} 2^{Loc}$ of the first- and second-order variables, respectively, we have $S, \iota, v \models_{mso} \text{Heap}(X) \iff v(X) = \text{dom}(h)$.*

Proof. By definition of $\text{Heap}(X)$, $v(X)$ is the set of locations $\ell \in Loc$ such that $h_s(\ell) \neq \perp$, for some $s = 1, \dots, \|Sel\|$. But this is exactly the definition of $\text{dom}(h)$, by Def. 1. \square

The following lemma says that an MSO formula obtained as a translation of a basic spatial SL formula is true in a state S if and only if it is true on any extension of S .

Lemma 9. *Let σ be a basic spatial SL formula, $S = \langle s, h \rangle$ be any state, $\iota : LVar_{mso} \rightarrow_{fin} Loc$, $v : LVar_{mso} \rightarrow_{fin} 2^{Loc}$ be interpretations of first and second-order variables, respectively. Then, for any state S' , such that $S \uplus S'$ is defined, we have:*

$$S, \iota, v[X \leftarrow \text{dom}(h)] \models_{mso} \overline{\sigma}(X) \iff S \uplus S', \iota, v[X \leftarrow \text{dom}(h)] \models_{mso} \overline{\sigma}(X)$$

Proof. By induction on the structure of σ . \square

Proof of Lemma 3: By induction on the structure of φ . The most interesting case is the separating conjunction, i.e. $\varphi \equiv \sigma_1 * \sigma_2$ for two spatial SL formulae σ_1 and σ_2 .

“ \Rightarrow ” $S, \iota \models_{sl} \sigma_1 * \sigma_2$ if and only if there exist two states $S_1 = \langle s, h_1 \rangle$ and $S_2 = \langle s, h_2 \rangle$ such that $S_i, \iota \models_{sl} \sigma_i$, for both $i = 1, 2$ and $S = S_1 \uplus S_2$. By the induction hypothesis we have that $S_i, \bar{\iota}, v[Y_i \leftarrow \text{dom}(h_i)] \models_{mso} \overline{\sigma}_i(Y_i) \wedge \text{Heap}(Y_i)$, and by Lemma 9, we have $S_1 \uplus S_2, \bar{\iota}, v[Y_i \leftarrow \text{dom}(h_i)] \models_{mso} \overline{\sigma}_i(Y_i)$, for both $i = 1, 2$. Hence:

$$\begin{aligned} S_1 \uplus S_2, \bar{\iota}, v[Y_1 \leftarrow \text{dom}(h_1)][Y_2 \leftarrow \text{dom}(h_2)] &\models_{mso} \overline{\sigma}_1(Y_1) \wedge \overline{\sigma}_2(Y_2) \\ S_1 \uplus S_2, \bar{\iota}, v[X \leftarrow \text{dom}(h_1) \cup \text{dom}(h_2)] &\models_{mso} \exists Y_1 \exists Y_2 . \overline{\sigma}_1(Y_1) \wedge \overline{\sigma}_2(Y_2) \wedge \Pi(Y_1, Y_2, X) \\ S_1 \uplus S_2, \bar{\iota}, v[X \leftarrow \text{dom}(h_1) \cup \text{dom}(h_2)] &\models_{mso} \overline{\varphi}(X) \end{aligned}$$

By Proposition 1, we obtain further:

$$S_1 \uplus S_2, \bar{\iota}, v[X \leftarrow \text{dom}(h_1) \cup \text{dom}(h_2)] \models_{mso} \text{Heap}(X)$$

hence $S, \bar{\iota}, v[X \leftarrow \text{dom}(h)] \models_{mso} \overline{\varphi}(X) \wedge \text{Heap}(X)$.

“ \Leftarrow ” If $S, \bar{\iota}, v[X \leftarrow \text{dom}(h)] \models_{mso} \exists Y_1 \exists Y_2 . \overline{\sigma}_1(Y_1) \wedge \overline{\sigma}_2(Y_2) \wedge \Pi(Y_1, Y_2, X) \wedge \text{Heap}(X)$, then there exists two sets of locations, call them L_1 and L_2 , such that $L_1 \cap L_2 = \emptyset$ and

$L_1 \cup L_2 = \text{dom}(h)$, such that $S, \bar{\tau}, v[Y_i \leftarrow L_i] \models_{mso} \overline{\sigma_i}(Y_i)$, for both $i = 1, 2$. Let h_1 and h_2 be the restrictions of h to L_1 and L_2 , respectively, and define $S_i = \langle s, h_i \rangle$, for both $i = 1, 2$. Clearly $S_1 \uplus S_2 = S$. By Lemma 9 we have $S_i, \bar{\tau}, v[Y_i \leftarrow \text{dom}(h_i)] \models_{mso} \overline{\sigma_i}(Y_i)$. By Proposition 1 we have, moreover that $S_i, \bar{\tau}, v[Y_i \leftarrow \text{dom}(h_i)] \models_{mso} \text{Heap}(Y_i)$, for both $i = 1, 2$. Applying the induction hypothesis, we obtain $S_i, \mathbf{1} \models_{sl} \sigma_i$, for both $i = 1, 2$, hence $S, \mathbf{1} \models_{sl} \varphi$. \square

Proof of Lemma 4: The proof relies on the following claim:

Claim. For any two positions $p, r \in \text{dom}(t)$ such that either (i) p is a child of r , (ii) r is a child of p , or (iii) $p = r$ we have:

$$x^p = y^r \text{ occurs in } \phi_t \iff A_\varphi \text{ moves in one step from } \langle p, q_x^{\text{var}} \rangle \text{ to } \langle r, q_y^{\text{var}} \rangle$$

Proof. We give the proof for the second case, the rest of the cases being similar. Assume that $r = p.k$, for some $k \in \mathcal{D}_+(\varphi)$. “ \Rightarrow ” By the definition of ϕ_t , $x^p = y^{p.k}$ occurs in ϕ_t^p only if $t(p) \equiv R_{ij}$, $(\text{tail}(R_{ij}))_k = P_{i_k}(y_1, \dots, y_{n_{i_k}})$, where $P_{i_k}(x_{i_k,1}, \dots, x_{i_k,n_{i_k}})$ is the corresponding definition in φ , $x \equiv x_{i_k,\ell}$, and $y \equiv y_\ell$, for some $\ell = 1, \dots, n_{i_k}$. In this case we have $(q_{x_{i_k},\ell}^{\text{var}}, k) \in \Delta(q_y^{\text{var}}, R_{ij}^s, \tau)$, for all $s \in \mathcal{D}(\varphi)$ and all $\tau \in \Sigma_\varphi \cup \{?\}$, and the conclusion follows. “ \Leftarrow ” By definition, A_φ has a transition rule $(q_y^{\text{var}}, k) \in \Delta(q_x^{\text{var}}, \sigma^s, \tau)$ for an $s \in \mathcal{D}(\varphi)$ only if $\sigma \equiv R_{ij}$, $(\text{tail}(R_{ij}))_k \equiv P_\ell(y_1, \dots, y_{n_\ell})$, $y \equiv x_{\ell,j}$ and $x \equiv y_j$, for some $j = 1, \dots, n_\ell$. In this case, the equality $x^p = y^r$ occurs in ϕ_t^p . \square

“ \Rightarrow ” $x^p = y^r$ is implied by ϕ_t only if there exists a path $p = p_1, \dots, p_n = r$ in $\text{dom}(t)$, and variables $x \equiv z_1, z_2, \dots, z_{n-1}, z_n \equiv y \in \text{LVar}_{sl}$, such that $z_i^{p_i} = z_{i+1}^{p_{i+1}}$ occurs in ϕ_t , for all $i = 1, \dots, n-1$. By the above claim, A_φ has a run from $\langle p, q_x^{\text{var}} \rangle$ to $\langle r, q_y^{\text{var}} \rangle$ along this path. “ \Leftarrow ” If A_φ has a run from $\langle p, q_x^{\text{var}} \rangle$ to $\langle r, q_y^{\text{var}} \rangle$ over t , there exist a sequence of positions $p = p_1, p_2, \dots, p_{n-1}, p_n = q \in \text{dom}(t)$ and variables $x \equiv z_1, z_2, \dots, z_{n-1}, z_n \equiv y$ such that A_φ moves in one step from $\langle p_i, q_{z_i}^{\text{var}} \rangle$ to $\langle p_{i+1}, q_{z_{i+1}}^{\text{var}} \rangle$. By the above claim, there exist equalities $z_i^{p_i} = z_{i+1}^{p_{i+1}}$ occurring in ϕ_t , hence $x^p = y^r$ is a consequence of ϕ_t . \square

Proof of Lemma 5: “ \Rightarrow ” If $S, \mathbf{1} \models_{sl} P_i(x_1, \dots, x_n)$ then $S, \mathbf{1} \models_{sl} \phi_t$ for some unfolding tree $t \in \mathcal{T}_i(\varphi)$. By induction on the structure of t , one can build a bijective tree $\mu_0 : \text{dom}(t) \rightarrow \text{dom}(h)$, and define sets $S_{ij}^d = \{\ell \in \text{dom}(h) \mid t(\mu_0^{-1}(\ell)) \equiv R_{ij}$ and $\exists p \in \text{dom}(t). \mu_0^{-1}(\ell) = p.d\}$ for all $d \in \mathcal{D}_+(\varphi)$ and $S_{ij}^{-1} = \{\ell \in \text{dom}(h) \mid \mu_0^{-1}(\ell) = \varepsilon \text{ and } t(\varepsilon) \equiv R_{ij}\}$. Let $\bar{\tau}' = \bar{\tau}[r \leftarrow \mu_0(\varepsilon)]$ and $v : \text{LVar}_{mso} \xrightarrow{\text{fin}} 2^{\text{Loc}}$ be any interpretation of second order variables such that $v(X_{ij}^d) = S_{ij}^d$ and $v(T) = \text{dom}(h)$. By Proposition 1, we have $S, \bar{\tau}', v \models_{mso} \text{Heap}(T)$. Next, we prove the following conditions:

1. $S, \bar{\tau}', v \models_{mso} \text{backbone}_i(r, \vec{\mathbf{X}}, T)$
2. $S, \bar{\tau}', v \models_{mso} \text{inner_edges}(r, \vec{\mathbf{X}}, T)$
3. $S, \bar{\tau}', v \models_{mso} \text{no_double_alloc}(r, \vec{\mathbf{X}}, T)$
4. $S, \bar{\tau}', v \models_{mso} \text{param}_{i,j}(r, \vec{\mathbf{X}}, T)$, for all $j = 1, \dots, n_i$

(1) Let $\lambda_0 : \text{dom}(t) \rightarrow \{X_1, \dots, X_m\}$ be a function defined as $\lambda_0(p) = X_j$ iff $\mu_0(p) \in S_j$. It is immediate that $v(X_i) = \{\ell \in \text{dom}(h) \mid \lambda_0(\mu_0^{-1}(\ell)) = X_i\}$. By Lemma 6 (Appendix A) it follows that:

$$S, \bar{\tau}', v \models_{mso} \text{tree}(r, \vec{\mathbf{X}}, T) \wedge X_i^{-1}(r)$$

To show:

$$S, \vec{t}', v \models_{mso} \text{succ_labels}(\vec{\mathbf{X}})$$

let X_{ij}^k be an arbitrary variable from $\vec{\mathbf{X}}$ and let $\ell \in S_{ij}^k$ be a location. Hence $t(\mu_0^{-1}(\ell)) \equiv R_{ij}$. Suppose that $\text{tail}(R_{ij}) = \langle P_{k_1}, \dots, P_{k_{r_{ij}}} \rangle$, and let $m = \mu_0(\mu_0^{-1}(\ell).d)$ for some arbitrary $d \in \{1, \dots, r_{ij}\}$. Clearly $m \in S_{k_d j}^d$, for some $j = 1, \dots, n_{k_d}$. One can now easily check that:

$$S, \vec{t}'[x \leftarrow \ell][y \leftarrow m], v \models_{mso} \text{succ}_d(x, y, \vec{\mathbf{X}})$$

which concludes this point.

(2) If $S, \vec{t}', v \models_{mso} \exists \vec{\mathbf{Y}} . \Phi_{A_P}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$ then there exist sets $U_1, \dots, U_k \subseteq \text{Loc}$ such that:

$$S, \vec{t}', v[\vec{\mathbf{Y}} \leftarrow \vec{\mathbf{U}}] \models_{mso} \Phi_{A_P}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}})$$

By Lemma 6, there exists a unique prefix-closed set $P \subseteq \mathbb{N}^*$ a unique bijective tree $\mu : P \rightarrow \text{dom}(h)$ and a unique tree $\lambda : P \rightarrow \{X_1, \dots, X_m\}$ meeting the three properties of Lemma 6. Since $\text{dom}(t)$, μ_0 and λ_0 meet the requirements of Lemma 6, it turns out that $P = \text{dom}(t)$, $\mu = \mu_0$ and $\lambda = \lambda_0$. By Lemma 8, A_P has an accepting run π over λ_0 , such that $U_j = v(Y_j) = \{\mu_0(p) \mid \langle p, q_j \rangle \text{ occurs on } \pi\}$. Let ℓ, m be two arbitrary locations such that:

$$S, \vec{t}'[x \leftarrow \ell][y \leftarrow m], v[\vec{\mathbf{Y}} \leftarrow \vec{\mathbf{U}}] \models_{mso} Y_s^{\text{sel}}(x) \wedge Y_f(y)$$

By Lemma 8, there exist positions $p_1, p_2 \in \text{dom}(t)$, such that $\mu_0(p_1) = \ell$ and $\mu_0(p_2) = m$, and variables $\alpha, \beta \in LVar_{sl}$ such that A_P has a run:

$$\langle p_1, q_s^{\text{sel}} \rangle, \langle p_1, q_\alpha^{\text{var}} \rangle \dots \langle p_2, q_\beta^{\text{var}} \rangle, \langle p_2, q_f \rangle$$

over λ_0 , and implicitly, over t . Notice that, by the definition of A_P , q_s^{sel} and q_f occur exactly once on each accepting run, and moreover, q_f is the final state on the run.

Hence, by Lemma 4, $S, \vec{t} \models_{sl} \alpha^{p_1} = \beta^{p_2}$. Since ℓ and m are allocated at p_1 and p_2 in t , respectively, there exists an edge $\ell \xrightarrow{s} m$ in S . We have, subsequently:

$$S, \vec{t}'[x \leftarrow \ell][y \leftarrow m], v \models_{mso} \text{edges}_s(x, y)$$

which concludes this point.

(3) By contradiction, let us suppose that there exist two distinct locations ℓ and m such that:

$$S, \vec{t}'[x \leftarrow \ell][y \leftarrow m], v[\vec{\mathbf{Y}} \leftarrow \vec{\mathbf{U}}] \models_{mso} \Phi_{B_P}(r, \vec{\mathbf{X}}, T, \vec{\mathbf{Y}}) \wedge Y_0(x) \wedge Y_f(y)$$

By an argument similar to the one from point (2), there exist two variables $\alpha, \beta \in LVar_{sl}$ such that B_P has a run:

$$\langle \mu_0^{-1}(\ell), q_0 \rangle, \langle \mu_0^{-1}(\ell), q_\alpha^{\text{var}} \rangle, \dots, \langle \mu_0^{-1}(m), q_\beta^{\text{var}} \rangle, \langle \mu_0^{-1}(m), q_f \rangle$$

over λ_0 , or equivalently, over t . By the definition of B_P , $\alpha^{\mu_0^{-1}(\ell)} = \beta^{\mu_0^{-1}(m)}$ is a consequence of ϕ_t and moreover both α and β are allocated at positions $\mu_0^{-1}(\ell)$ and $\mu_0^{-1}(m)$

in t , respectively. The latter facts contradict with the hypothesis that $S, \mathbf{i} \models_{sl} \phi_t$, since, in this case, ϕ_t would not be satisfiable, according to the semantics of SL.

(4) This point follows the case split in the definition of $param_{i,j}$ and is proved among the same lines as point (2) above.

“ \Leftarrow ” If $S, \bar{\mathbf{i}}, \mathbf{v}[T \leftarrow dom(h)] \models_{mso} \bar{P}_i(\bar{x}_1, \dots, \bar{x}_k, T)$, then there exists a location $\ell \in Loc$ and sets $L_1, \dots, L_m \subseteq Loc$ such that:

- i. $S, \bar{\mathbf{i}}[r \leftarrow \ell], \mathbf{v}[\vec{\mathbf{X}} \leftarrow \vec{\mathbf{L}}][T \leftarrow dom(h)] \models_{mso} backbone_i(r, \vec{\mathbf{X}}, T)$
- ii. $S, \bar{\mathbf{i}}[r \leftarrow \ell], \mathbf{v}[\vec{\mathbf{X}} \leftarrow \vec{\mathbf{L}}][T \leftarrow dom(h)] \models_{mso} inner_edges(r, \vec{\mathbf{X}}, T)$
- iii. $S, \bar{\mathbf{i}}[r \leftarrow \ell], \mathbf{v}[\vec{\mathbf{X}} \leftarrow \vec{\mathbf{L}}][T \leftarrow dom(h)] \models_{mso} no_double_alloc(r, \vec{\mathbf{X}}, T)$
- iv. $S, \bar{\mathbf{i}}[r \leftarrow \ell], \mathbf{v}[\vec{\mathbf{X}} \leftarrow \vec{\mathbf{L}}][T \leftarrow dom(h)] \models_{mso} param_{i,j}(r, \vec{\mathbf{X}}, T)$, for all $1 \leq i \leq n$ and all $1 \leq j \leq n_i$

For simplicity, we denote $\bar{\mathbf{i}}' = \mathbf{i}[r \leftarrow \ell]$ and $\mathbf{v}' = \mathbf{v}[\vec{\mathbf{X}} \leftarrow \vec{\mathbf{L}}][T \leftarrow dom(h)]$ in the rest of this proof. By (i) and Lemma 6, there exist a set $P \subseteq \mathbb{N}^*$, a bijective tree $\mu: P \rightarrow dom(h)$, and a tree $\lambda: P \rightarrow \vec{\mathbf{X}}$ such that $L_i = \{\ell \in dom(h) \mid \lambda(\mu^{-1}(\ell)) = X_i\}$. Since each variable X_{ij} from $\vec{\mathbf{X}}$ corresponds one-to-one to the rule R_{ij} from \mathcal{P} , we can build a tree $t: P \rightarrow \vec{\mathbf{R}}$ as $t(p) \stackrel{def}{=} R_{ij}$ iff $\lambda(p) = X_{ij}$, for all $p \in P$. Since, by (i):

$$S, \bar{\mathbf{i}}', \mathbf{v}' \models_{mso} succ_labels(\vec{\mathbf{X}}) \wedge X_i^{-1}(r)$$

we obtain that t is an unfolding tree (Def. 3), and moreover $t \in \mathcal{T}_i(\mathcal{P})$. It remains to be shown that $S, \mathbf{i} \models_{sl} \phi_t$. To this end, we extend \mathbf{i} to an assignment $\mathbf{i}_t: LVar_{sl}^P \rightarrow Loc$ such that:

1. $\mathbf{i}_t(x^e) = \mathbf{i}(x)$, for all $x \in LVar_{sl}$
2. $\mathbf{i}_t(x^{p_1}) = \mathbf{i}(y^{p_2})$, for all $p_1, p_2 \in P$ and $x, y \in LVar_{sl}$ such that $\phi_t \rightarrow x^{p_1} = y^{p_2}$
3. for all $p \in P$, if $head(t(p)) \equiv x \mapsto (y_1, \dots, y_s)$ then $\mathbf{i}_t(x^p) \xrightarrow{i} \mathbf{i}_t(y_i^p)$ in S , for all $i = 1, \dots, s$, and moreover, there are no other outgoing edges from $\mathbf{i}_t(x^p)$ in S

Given $\bar{\mathbf{i}}, \mu$ and t , defined above, we define \mathbf{i}_t as follows:

- $\mathbf{i}_t(x_j^e) \stackrel{def}{=} \bar{\mathbf{i}}(\bar{x}_j)$, for all $j = 1, \dots, k$
- $\mathbf{i}_t(x^p) \stackrel{def}{=} \mu(p)$, for all $p \in P$ such that $head(t(p)) \equiv x \mapsto (\dots)$
- for all $x^p \in LVar^P$ not assigned previously, $\mathbf{i}_t(x^p) \stackrel{def}{=} \mathbf{i}_t(y^q)$ if and only if $\phi_t \rightarrow x^p = y^q$ and y^q is assigned by one of the above points

Since \mathcal{P} is established, every existentially quantified variable $x^p \in LVar^P$ that occurs in ϕ_t is connected to an allocated variable y^q , i.e. $head(t(q)) \equiv y^q \mapsto (\dots)$, by a path of equalities $x^p = z_1^{p_1} = \dots = z_n^{p_n} = y^q$ all occurring in ϕ_t . Hence \mathbf{i}_t assigns locations to all existentially quantified variables in ϕ_t . Clearly, \mathbf{i}_t satisfies points (1) and (2) above. To show that \mathbf{i}_t meets point (3), fix an arbitrary position $p_0 \in P$ such that $head(t(p_0)) \equiv x \mapsto (y_1, \dots, y_s)$. Observe first that, by (i), $\mathbf{i}_t(x^{p_0})$ has no outgoing edges $\mathbf{i}_t(x^{p_0}) \xrightarrow{i} \ell$, for any $i > s$. For the rest, let us fix some arbitrary $1 \leq i_0 \leq s$ and show that $\mathbf{i}_t(x^{p_0}) \xrightarrow{i_0} \mathbf{i}_t(y_{i_0}^{p_0})$ is an edge in S . There are two cases:

- y_k is an existentially quantified variable of ϕ_t
- y_k is a parameter x_j of the predicate P_i

We shall carry out the proof only in the first case, the reasoning being similar in the second. As previously discussed, if y_k is existentially quantified, there exists a sequence of equalities $y_{i_0}^{p_0} = z_1^{p_1} = \dots = z_n^{p_n}$ occurring in ϕ_t , such that z_n is allocated by $\text{head}(t(p_n))$. By Lemma 4, $A_{\mathcal{P}}$ has a run from $\langle p_0, q_{y_{i_0}}^{\text{var}} \rangle$ to $\langle p_n, q_{z_n}^{\text{var}} \rangle$ over t . Hence $A_{\mathcal{P}}$ has also a run:

$$\langle p_0, q_s^{\text{sel}} \rangle, \langle p_0, q_{y_k}^{\text{var}} \rangle, \dots, \langle p_n, q_{z_n}^{\text{var}} \rangle, \langle p_n, q_f \rangle$$

over t , and, equivalently, a loop-free run π over λ . Let $U_j = \{\mu(p) \mid \langle p, q_j \rangle \text{ occurs on } \pi\}$. By Lemma 8, we obtain:

$$S, \vec{v}'[\vec{Y} \leftarrow \vec{U}] \models_{\text{mso}} \Phi_{A_{\mathcal{P}}}(r, \vec{X}, T, \vec{Y})$$

and moreover, $\mu(p_0) \in U_s^{\text{sel}}$ and $\mu(p_n) \in U_f$, where U_s^{sel} and U_f are the sets of locations corresponding to the states q_s^{sel} and q_f , respectively. We obtain, further:

$$S, \vec{v}'[x \leftarrow \mu(p_0)][y \leftarrow \mu(p_n)], v' \models_{\text{mso}} \exists \vec{Y}. \Phi_{A_{\mathcal{P}}}(r, \vec{X}, T, \vec{Y}) \wedge Y_s^{\text{sel}}(x) \wedge Y_f(y)$$

By (ii), we obtain:

$$S, \vec{v}'[x \leftarrow \mu(p_0)][y \leftarrow \mu(p_n)], v' \models_{\text{mso}} \text{edges}(x, y)$$

hence the conclusion follows. \square

Lemma 10. *Let $P_i(x_{i,1}, \dots, x_{i,n})$ be a predicate of a recursive definition system \mathcal{P} , $S = \langle s, h \rangle$ be a state, and $\mathbf{v} : L\text{Var}_{\text{mso}} \rightarrow_{\text{fin}} \text{Loc}$ and $\mathbf{v}' : L\text{Var}_{\text{mso}} \rightarrow_{\text{fin}} 2^{\text{Loc}}$ be interpretations of first and second-order variables, respectively. Then, for any state S' , such that $S \uplus S'$ is defined, we have:*

$$\begin{aligned} S, \mathbf{v}[T \leftarrow \text{dom}(h)] &\models_{\text{mso}} \overline{P}_i(\overline{x_{i,1}}, \dots, \overline{x_{i,n}}, T) \\ \iff S \uplus S', \mathbf{v}[T \leftarrow \text{dom}(h)] &\models_{\text{mso}} \overline{P}_i(\overline{x_{i,1}}, \dots, \overline{x_{i,n}}, T) \end{aligned}$$

Proof. The proof is done by inspection of $\overline{P}(\overline{x_1}, \dots, \overline{x_n}, T)$. Namely we need to prove the following equivalences, for some $\ell \in \text{Loc}$ and sets S_1, \dots, S_m corresponding to the variables X_1, \dots, X_m :

$$\begin{aligned} S, \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{backbone}_i(r, \vec{X}, T) &\iff S \uplus S', \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{backbone}_i(r, \vec{X}, T) \\ S, \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{inner_edges}(r, \vec{X}, T) &\iff S \uplus S', \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{inner_edges}(r, \vec{X}, T) \\ S, \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{no_double_alloc}(r, \vec{X}, T) &\iff S \uplus S', \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{no_double_alloc}(r, \vec{X}, T) \\ S, \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{param}_{ij}(r, \vec{X}, T) &\iff S \uplus S', \mathbf{v}', \mathbf{v}' \models_{\text{mso}} \text{param}_{ij}(r, \vec{X}, T) \end{aligned}$$

for all $1 \leq j \leq n$, where $\mathbf{v}' = \mathbf{v}[r \leftarrow \ell]$ and $\mathbf{v}' = \mathbf{v}[\vec{X} \leftarrow \vec{S}][T \leftarrow \text{dom}(h)]$. These equivalences can be proved by case analysis. \square

Proof of Theorem 3: Let us first consider the case $k = 0$, i.e. ϕ is a basic SL formula $\phi \equiv \exists \mathbf{z} . \phi(\mathbf{y}_0)$. By Lemma 3, for any state $S = \langle s, h \rangle$, we have that $S \models_{\text{sl}} \phi$ if and only

if $S, \iota \models_{mso} \overline{\Phi}(X) \wedge \text{Heap}(X)$, where ι is any interpretation such that $\iota(X) = \text{dom}(h)$. Hence $S \models_{mso} \exists X . \overline{\Phi}(X) \wedge \text{Heap}(X)$. Dually, if $S \models_{mso} \exists X . \overline{\Phi}(X) \wedge \text{Heap}(X)$, then $S, \iota[X \leftarrow L] \models_{mso} \overline{\Phi}(X) \wedge \text{Heap}(X)$, where $L \subseteq \text{Loc}$ is a set of locations. By Proposition 1, we have $L = \text{dom}(h)$. Hence $S, \iota \models_{sl} \overline{\Phi}$, by Lemma 3.

The case $k > 0$ is dealt with by induction on k . For $k = 1$ we have $\overline{\Phi} \equiv \exists \mathbf{z} . \overline{\phi}(\mathbf{y}_0) * P(\mathbf{y}_1)$. “ \Rightarrow ” If $S \models_{sl} \exists \mathbf{z} . \overline{\phi}(\mathbf{y}_0) * P(\mathbf{y}_1)$ then $S, \iota \models_{sl} \overline{\phi}(\mathbf{y}_0) * P(\mathbf{y}_1)$ for some interpretation $\iota : LVar_{sl} \rightarrow_{fin} \text{Loc}$, such that $\iota(x) \neq \perp$, for all $x \in \mathbf{z}$. Hence there exists two states $S_1 = \langle s, h_1 \rangle$ and $S_2 = \langle s, h_2 \rangle$ such that $S_1 \uplus S_2 = S$, and moreover $S_1, \iota \models_{sl} \overline{\phi}(\mathbf{y}_0)$ and $S_2, \iota \models_{sl} P(\mathbf{y}_1)$. Applying Lemma 3 and 5, respectively, we obtain:

$$\begin{aligned} S_1, \bar{\iota}[X_0 \leftarrow \text{dom}(h_1)] &\models_{mso} \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \wedge \text{Heap}(X_0) \\ S_2, \bar{\iota}[X_1 \leftarrow \text{dom}(h_2)] &\models_{mso} \overline{P}(\bar{\mathbf{y}}_1, X_1) \wedge \text{Heap}(X_1) \end{aligned}$$

where $\bar{\iota} : LVar_{mso} \rightarrow_{fin} (\text{Loc} \cup 2^{\text{Loc}})$ is an interpretation meeting the requirements of both Lemma 3 and 5. Applying Lemma 9 and 10, respectively, we obtain:

$$\begin{aligned} S, \bar{\iota}[X_0 \leftarrow \text{dom}(h_1)] &\models_{mso} \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \\ S, \bar{\iota}[X_1 \leftarrow \text{dom}(h_2)] &\models_{mso} \overline{P}(\bar{\mathbf{y}}_1, X_1) \end{aligned}$$

Since $S_1 \uplus S_2 = S$, we obtain:

$$\begin{aligned} S, \bar{\iota}[X \leftarrow \text{dom}(h)] &\models_{mso} \exists X_0 \exists X_1 . \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \wedge \overline{P}(\bar{\mathbf{y}}_1, X_1) \wedge \Pi(X_0, X_1, X) \\ S, \bar{\iota}[X \leftarrow \text{dom}(h)] &\models_{mso} \exists \mathbf{z} \exists X_0 \exists X_1 . \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \wedge \overline{P}(\bar{\mathbf{y}}_1, X_1) \wedge \Pi(X_0, X_1, X) \end{aligned}$$

and by Proposition 1 we also have $S, \bar{\iota}[X \leftarrow \text{dom}(h)] \models_{mso} \text{Heap}(X)$. The conclusion follows. “ \Leftarrow ” If $S \models_{mso} \exists X . \overline{\Phi}(X) \wedge \text{Heap}(X)$, then for any interpretation $\bar{\iota} : LVar_{mso} \rightarrow_{fin} (\text{Loc} \cup 2^{\text{Loc}})$ we have $S, \bar{\iota}[X \leftarrow \text{dom}(h)] \models_{mso} \overline{\Phi}(X)$, by Proposition 1. Hence there exists $L_1, L_2 \subseteq \text{dom}(h)$, such that $L_1 \cap L_2 = \emptyset$, $L_1 \cup L_2 = \text{dom}(h)$, and:

$$\begin{aligned} S, \bar{\iota}[X_0 \leftarrow L_1] &\models_{mso} \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \\ S, \bar{\iota}[X_1 \leftarrow L_2] &\models_{mso} \overline{P}(\bar{\mathbf{y}}_1, X_1) \end{aligned}$$

Let h_1, h_2 be the restrictions of h to L_1, L_2 , respectively, and $S_1 = \langle s, h_1 \rangle$, $S_2 = \langle s, h_2 \rangle$. Clearly $S = S_1 \uplus S_2$. By Lemma 9 and 10, respectively, we have that:

$$\begin{aligned} S_1, \bar{\iota}[X_0 \leftarrow L_1] &\models_{mso} \overline{\phi}(\bar{\mathbf{y}}_0, X_0) \\ S_2, \bar{\iota}[X_1 \leftarrow L_2] &\models_{mso} \overline{P}(\bar{\mathbf{y}}_1, X_1) \end{aligned}$$

and by Lemma 3 and 5, respectively, we obtain:

$$\begin{aligned} S_1, \iota \models_{sl} \overline{\phi}(\mathbf{y}_0) \\ S_2, \iota \models_{sl} P(\mathbf{y}_1) \end{aligned}$$

for an interpretation $\iota : LVar_{sl} \rightarrow_{fin} \text{Loc}$ meeting the conditions of Lemma 3 and 5. Hence $S, \iota \models_{sl} \overline{\phi}(\mathbf{y}_0) * P(\mathbf{y}_1)$, which leads to $S \models_{sl} \exists \mathbf{z} . \overline{\phi}(\mathbf{y}_0) * P(\mathbf{y}_1)$.

The induction step follows a similar argument. \square